arXiv:1406.3790v1 [cs.OH] 15 Jun 2014

High Performance Network-on-Chips (NoCs) Design: Performance Modeling, Routing Algorithm and Architecture Optimization

by

Zhiliang Qian

A Thesis Submitted to
The Hong Kong University of Science and Technology
in Partial Fulfillment of the Requirements for
the Degree of Doctor of Philosophy
in the Department of Electronic and Computer Engineering

March 2014, Hong Kong

Authorization

I hereby declare that I am the sole author of the thesis.

I authorize the Hong Kong University of Science and Technology to lend this thesis to other institutions or individuals for the purpose of scholarly research.

I further authorize the Hong Kong University of Science and Technology to reproduce the thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Zhiliang Qian

March 2014

High Performance Network-on-Chips (NoCs) Design: Performance Modeling, Routing Algorithm and Architecture Optimization

by

Zhiliang Qian

This is to certify that I have examined the above Ph.D. thesis
and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by
the thesis examination committee have been made.

Prof. Chi-Ying TSUI (ECE), Thesis Supervisor

Prof. Albert C S CHUNG (CSE), Thesis Committee Chairman

Prof. Chin-Tau LEA (ECE), Thesis Committee Member

Prof. Jiang XU (ECE), Thesis Committee Member

Prof. Chi-Keung TANG (CSE), Thesis Committee Member

Prof. Ross D. MURCH, Head of Department

Department of Electronic and Computer Engineering

March 2014

ACKNOWLEDGMENTS

This thesis summarizes my five and half years study and research experience in VLSI Research Laboratory at Hong Kong University of Science and Technology. I would like to take this opportunity to thank all people who have helped, accompanied and supported me during my PhD study in HKUST.

My foremost thanks belong to my supervisor, Prof. Chi-Ying Tsui, for his patient guidance, encouragement, understanding, and support all over the time. His brilliant insights on VLSI and highly motivation on NoC led me to the wonderful world of on-chip networks. He always inspires me and has provided me with valuable advice in my study. His brilliance, enthusiasm and hard working towards research, and nice guidance not only have benefited my current research but also have a lasting influence in my professional career and personal development.

I am deeply grateful to Prof. Chin Tau Lea, Prof. Jiang Xu, Prof. Chi-Keung Tang and Prof. Albert C. S. Chung for serving on my Thesis Committee, and I am full of gratitude to Prof. Oliver Chiu-Sing CHOY from Chinese University of Hong Kong for being my Thesis External Examiner. They managed to take time to serve on my committee in their tight schedule. I would also like to thank Prof. Radu Marculescu from Carnegie Mellon University for hosting my Fulbright visit. I benefit a lot through working together with Prof. Radu Marculescu, Prof. Diana Marculescu in CMU and Prof. Paul Bogdan from University of Southern California. Their valuable feedback helped improve my work in many ways.

I would like to thank Mr. Leo Fok and Mr Siu Fai Luk for providing me with great technical and administrative support these years. I thank everyone in VLSI lab (especially Dr. Jie Jin, Dr. Liu Feng, Dr. Shao Hui, Mr Yunxiao ling, Mr Yingfei Teh, Mr Xing Li, Mr Youzhe Fan, Mr Jingyang Zhu and Mr Syed Abbas Mohsin) for creating a family environment and the friends I met in HKUST for their friendship and help.

Last, but not the least, I would like to express my deepest gratitude to my parents and girlfriend for their unconditional love and support through all these years during my Ph.D. study.

TABLE OF CONTENTS

Title Page i

Authorization Page ii

Signature Page iii

Acknowledgments iv

Table of Contents v

List of Figures ix

List of Tables xiii

Abstract xiv

Chapter 1 Introduction 1

1.1 Challenges in computing platform design 1

1.2 NoCs for multi-core communication 2

1.3 Application-specific NoC design flow 7

1.4 NoC characterization 9

1.4.1 NoC topology 9

1.4.2 NoC switching technique 10

1.4.3 NoC router design 11

1.5 Thesis overview 13

Chapter 2 SVR-NoC: A Learning Based NoC Latency Performance Model 16

2.1 Introduction 16

2.2 Background 19

2.3 Proposed NoC analytical model 22

2.3.1 Basic assumptions and notations 22

2.3.2 End-to-end delay formulation 23

2.3.3 Link dependency analysis 24

2.3.4 GE-type traffic modeling 25

2.3.5 WH router modeling 26
2.3.6 Discussion on extensions to VC router modeling 30
2.4 SVR-NoC latency model 31
2.4.1 Channel and source queuing regression models 32
2.4.2 Overall SVR-NoC methodology 32
2.4.3 Feature extraction for training data 35
2.4.4 Support vector regression for f_{CQ} and f_{SQ} 38
2.5 Experimental results 42
2.5.1 Experimental setup and training data preparation 42
2.5.2 Proposed queuing model accuracy 42
2.5.3 SVR prediction accuracy 43
2.5.4 Runtime comparison 45
2.6 Conclusion 51

Chapter 3 Thermal-aware and Application-specific Routing Algorithm Design 52

3.1 Introduction 52
3.2 Related Work 53
3.3 Methodology overview 55
3.3.1 Assumptions and preliminaries 55
3.3.2 Motivation for thermal-aware routing 58
3.3.3 Thermal-aware routing design flow 59
3.4 Main algorithm 60
3.4.1 Application-specific path set finding algorithm 61
3.4.2 Optimal traffic allocation 63
3.5 Router microarchitecture 69
3.6 Experimental results 70
3.6.1 Simulation environment setup 70
3.6.2 Adaptivity improvement 71
3.6.3 Peak energy simulation results 71
3.6.4 Simulation results with different processor/router energy ratio 76
3.7 Conclusion 77

Chapter 4 Fault-tolerant Routing Scheme for NoCs using Dynamic Reconfiguration of Partial-Faulty Routing Resources 80

4.1 Introduction 80
4.2 Background 82
4.3 Fault-tolerant NoC design 83
4.3.1 Fault-diagnosis and hybrid fault model 83
4.3.2 Dynamic buffer swapping (DBS) 85
4.3.3 Dynamic MUX swapping (DMS) 89
4.3.4 Handling hard link faults 89
4.3.5 Deadlock issue and handling strategy 90
4.4 Simulation Results 92
4.4.1 Simulation environment setup 92
4.4.2 Comparison of connectivity of the network 92
4.4.3 Improvement in packet acceptance rate 93
4.4.4 Effects on latency performance 94
4.5 Conclusion 96

Chapter 5 FSNoC: A Flit-level Speedup Scheme For NoCs Using Self-Reconfigurable Bi-directional Channels 97

5.1 Introduction 97
5.2 Motivations of FSNoC 101
5.2.1 A motivational example 101
5.2.2 Single router throughput comparisons 106
5.3 Implementation of FSNoC 108
5.3.1 Inter-router channel direction control 108
5.3.2 Router datapath design of FSNoC 114
5.4 FSNoC for long-wire links 117
5.4.1 Long-wire links with repeaters 118
5.4.2 Long-wire links with pipeline registers 120
5.5 Experimental results 122
5.5.1 Simulation setup 122
5.5.2 Simulation results on synthetic traffics 123
5.5.3 Simulation results on real benchmarks 128
5.5.4 Simulation results for FSNoC with link pipelines 129
5.5.5 Implementation overhead 131
5.6 Conclusion 132

Chapter 6 A Traffic-aware Adaptive Routing Algorithm on a Highly Reconfigurable Network-on-Chip Architecture **134**

6.1 Introduction 134

6.2 Background 136

6.3 New NoC architecture 138

6.3.1 Motivation for adding hub routers 138

6.3.2 Overview of proposed NoC architecture 139

6.3.3 Router implementation 140

6.4 Routing algorithm design 143

6.4.1 Choosing the express hub and EVC paths 143

6.4.2 Adaptive routing for BiNoC 144

6.5 Simulation results 150

6.5.1 Simulation setup 150

6.5.2 Comparison of adaptive routing for BiNoC 150

6.5.3 Effect of adding EVC and hub routers 152

6.5.4 Results using real world workloads 153

6.5.5 Overhead evaluation 155

6.6 Conclusion 156

Chapter 7 Conclusions and future work **157**

7.1 Research summary 157

7.2 Future research directions 158

References **160**

LIST OF FIGURES

1.1 The plots of transistor counts against dates and Moore's Law (courtesy of [1]) 3

1.2 The illustration of the processor performance growth since the late 1970s (from [2]) 3

1.3 Power density trends in Intel's CPU (from [3]) 4

1.4 Growth in clock rate of microprocessors (from [2]) 4

1.5 Manycore system roadmap for improving computing platform performance (from [4]) 5

1.6 The trend of on-chip interconnections (from [5]) 7

1.7 The NoC-based Video Object Plane Decoder (VOPD) platform in [6] 9

1.8 NoC synthesis flow for application specific multicore systems 10

1.9 The NoC router architecture with virtual channels and the open-loop source queuing delay measurement 12

2.1 The NoC performance model used in the synthesis inner loop for large design space exploration 17

2.2 An example showing the flow delay: a) the core communication graph, b) the links involved to calculate the end-to-end delay of flow $f_{0,8}$ c) the acquisition time h and transfer time η for link $l_{1.2}$ in WH model, d) h and η in VC model 21

2.3 Modeling of the GE packet generation process in [7] 26

2.4 An illustration of the channel service time of link l_i under WH routing 26

2.5 The transition diagram to calculate the average degree of virtual channel multiplexing [7] 31

2.6 SVR-NoC methodology overview 34

2.7 Illustration of the channel and source queuing feature vectors 34

2.8 The soft margin loss setting in ϵ-SVR [8] 38

2.9 An example RMSE surface in the grid search 41

2.10 Analytical model accuracy comparisons 43

2.11 a) Regression curve of the channel queuing time (CQ) b)Comparisons of the learning and queuing model 44

2.12 Comparisons with the queuing model under different traffic pattern and NoC configurations 44

2.13 Comparisons with the queuing model under different traffic pattern and NoC configurations -2 45

2.14 Latency prediction performance under trained traffic patterns 46
2.15 Latency prediction for trained synthetic traffics -2 47
2.16 Latency prediction for trained synthetic traffics -3 48
2.17 Latency prediction for untrained synthetic traffics 49
2.18 Latency prediction for untrained synthetic traffics -2 50
2.19 Flow latency comparison for DVOPD benchmark 50
2.20 Run time comparisons of the learning model and simulation 51
3.1 An example of: a) Core communication graph (CCG), b) NoC architecture, c) Channel dependency graph (CDG) and its strongly connected components(SCC) after mapping 56
3.2 A motivation example of allocating traffic among paths: a) the core communication graph(CCG), b) routing paths allocation on 2×3 NoC, c) two strategies of using the routing paths, d) thermal profile comparison (left: strategy one, right: strategy two) 58
3.3 Proposed design flow for the thermal-aware routing 60
3.4 Application specific and deadlock free path finding algorithm 61
3.5 Proof of deadlock free inheritance 67
3.6 Block diagram of the router supporting ratio-based routing: a)Router microarchitecture, b)Routing computation unit, c)Ratio-based output port selection 69
3.7 Adaptivity comparisons for different benchmarks 71
3.8 Peak energy simulation for real benchmarks 72
3.9 Latency and peak energy simulation for a) Random, b) Transpose-1 and c) Transpose-2 traffic 73
3.10 Latency and peak energy simulation for a) Hotspot-Rs, b) Hotspot-Tr and c) Bursty traffic 74
3.11 Latency and peak energy simulation for a) Hotspot-C and b) Butterfly traffic 75
3.12 An example of the NoC energy profile under hotspot-4c traffic: a) proposed routing b) oddeven routing c) negativefirst routing d) XY routing 76
4.1 A Comparison of various fault model 82
4.2 Architecture of fault-tolerant NoC [9, 10] 84
4.3 Dynamic Buffer Swapping (DBS) technique 87
4.4 A pathological case for packets interleaving 88
4.5 DMS hardware implementation 90
4.6 Deadlock situation for DBS operation 91
4.7 Connectivity comparison of the hybrid fault model 93

4.8 Comparison between buffer swapping strategy 95

4.9 Latency comparison of various fault model and schemes 95

5.1 Typical four-stage pipelined router micro-architecture using VC-based flow-control 98

5.2 The comparisons of different architectures a) example traffic flows b) timing diagram under typical uni-directional NoC c) timing diagram for typical NoC with 2X internal bandwidth d) timing diagram for BiNoCs 101

5.3 The proposed flit-level speedup switching timing diagram for the case in Fig. 5.2-a 103

5.4 Profiling results on an 8×8 mesh under different traffic patterns: a) Histogram of percentage of times that no flows in the opposite link direction at run time b) Breakdown of scenarios III and IV under the condition that there is no contention in the opposite link 104

5.5 The histogram of β on an 8×8 mesh under 80% saturation injection rate 106

5.6 Single router throughput comparisons of different router architectures: a) Throughput versus packet arrival rate $\bar{\beta}$ under $\bar{\alpha} = 1.0$. b) Throughput versus packet arrival rate $\bar{\beta}$ given $\bar{\alpha} = 0.5$ c) The throughput comparisons of four NoC architectures under a variety of $(\bar{\alpha}, \bar{\beta})$ combinations. 107

5.7 a) The channel direction control (CDC) module b) The timing diagram of link direction switching under the arrival of three packets (A,B,C) at router R1 and R2 109

5.8 Proposed buffer organization 115

5.9 Switch allocator logic (modified from [106]) 116

5.10 FSNoC design for long wire link with repeaters 118

5.11 a) FSNoC design for long wire links with pipeline registers, b) the detail schematics for data link pipeline registers, c) two types of conflicts during the the link traversal (LT) stage 119

5.12 Simulation results for six synthetic traffics on an 8×8 mesh under XY routing 124

5.13 The impacts of buffer depths and packet size on overall throughput 125

5.14 Latency comparisons for a) four architectures under oddeven (OE) routing algorithm b) FSNoC and BiNoC with WH router architecture and VC router architecture with different VC numbers c) NoC architectures with 4L links 126

5.15 Benchmark throughput comparison 127

5.16 Latency comparisons for the SPEC-web benchmarks 128

5.17 Histogram of the delivery time for Telecom 129

5.18 Latency comparisons for NoCs with 2-stage pipelined links 130

6.1 a) Proposed NoC architecture for an 8×8 mesh with bi-directional links, EVCs and hub routers b) The hub router architecture with each VC corresponding to a node in the region c) An example of routing from (1,3) to (3,1) using hub routers 136

6.2 Router architecture of a) BiNoC; b) BiNoC + EVC bypass node; c) BiNoC+EVC source/sink node 140

6.3 The express paths to the hub router 141

6.4 a)Router architecture of node (1,2) b) Contention between EVC and Hub paths 142

6.5 K-step (K=3) speculative routing 147

6.6 Comparison of the fitness functions for 8 × 8 meshes using bi-directional channels 149

6.7 Comparison of different adaptive routing schemes for 8 × 8 meshes 149

6.8 Comparison of different NoC architectures under various mesh size and traffic patterns. 151

6.9 Latency comparison for benchmarks 153

6.10 Percentage of packets using EVCs and hub routers 153

6.11 Energy comparison for benchmarks 154

6.12 Area overhead comparison 156

LIST OF TABLES

2.1 Summary and comparison of NoC latency models 17

2.2 Parameters and notations in NoC latency model 20

2.3 Training features for regression 36

3.1 Notations of application specific path set finding algorithm 62

3.2 Parameters and notations in LP formulation 64

3.3 Execution time for different mesh sizes 76

3.4 Peak energy reduction under various r_e(for Random, Transpose-1, Transpose-2 and Hotspot-C traffic) 77

3.5 Peak energy reduction under various r_e(for Butterfly, Hotspot-Tr,Hotspot-Rs and Bursty traffic) 78

3.6 Peak energy reduction under various r_e (for real benchmark traffic) 79

4.1 The proposed hybrid model for permanent faults 85

4.2 Handshake flow control for DBS 89

4.3 Packet acceptance rate comparison 94

4.4 Latency and packet acceptance rate comparison for various traffics 96

5.1 NoC architectures used in experiments 122

5.2 Area breakdown of different NoC architectures 132

5.3 Power breakdown of different NoC architectures 132

6.1 Notations in the routing for proposed NoC 145

High Performance Network-on-Chips (NoCs) Design: Performance Modeling, Routing Algorithm and Architecture Optimization

by

Zhiliang Qian

Department of Electronic and Computer Engineering

The Hong Kong University of Science and Technology

ABSTRACT

With technology scaling down, hundreds and thousands processing elements (PEs) can be integrated on a single chip. Consequently, the embedded systems have led to the advent of multi-core System-on-Chip (MPSoC) design and the high performance computer architectures have evolved into Chip Multi-processor (CMP) platforms. A scalable and modular solution to the interconnecting problem becomes critically important. Network-on-chip (NoC) has been proposed as an efficient solution to handle this distinctive challenge by providing efficient and scalable communication infrastructures among the on-chip resources. In this thesis, we have explored the high performance NoC design for MPSoC and CMP structures from the performance modeling in the offline design phase to the routing algorithm and NoC architecture optimization. More specifically, we first deal with the issue of how to estimate an NoC design fast and accurately in the synthesis inner loop. The simulation based evaluation method besides being slow, provides little insight to search the large design space in the NoC synthesis loop. Therefore, fast and accurate analytical models for NoC-based multicore performance evaluation are strongly desired to better explore the design space. For this purpose, we propose a machine learning based latency regression model to

evaluate the NoC designs with respect to different configurations before the system is built or taped-out. Then, for high performance NoC designs, we tackle one of the most important problems, *i.e.,* the routing algorithms design with different design constraints and objectives. For avoiding temperature hotspots, a thermal-aware routing algorithm is proposed to achieve an even temperature profile for application-specific Network-on-chips (NoCs). For improving the reliability, a routing algorithm to achieve maximum performance under fault is proposed. Finally, in the architecture level, we propose two new NoC structures using bi-directional links for the performance optimization. In particular, we propose a flit-level speedup scheme to enhance the network-on-chip(NoC) performance utilizing bidirectional channels. In addition to the traditional efforts on allowing flits of different packets using the idling internal and external bandwidth of the bi-directional channel, our proposed flit-level speedup scheme also allows flits within the same packet to be transmitted simultaneously on the bi-directional channel. We also propose a flexible NoC architecture which takes advantage of a dynamic distributed routing algorithm and improves the NoC communication performance with minimal energy overhead. This proposed NoC architecture exploits the self-reconfigurable bidirectional channels to increase the effective bandwidth and uses express virtual paths, as well as localized hub routers, to bypass some intermediate nodes at run time in the network. From the simulation results on both synthetic traffic and real workload traces, significantly performance improvement in terms of latency and throughput can be achieved.

CHAPTER 1

INTRODUCTION

1.1 Challenges in computing platform design

Computer and IC technology have made dramatic progress in the past few decades since the first generation of electronic computer was built [2]. As indicated by the Moore's Law [11] (shown in Fig. 1.1), the transistor count on an integrated circuits(ICs) doubles every two years. Accordingly, by the year of 2012, the state-of-the-art processors already contain billions of transistors (such as the Intel's 10-core Xeon CPU with 2.5 billion transistors [12] and Nvidia's 7.08 billion transistors GPU [13]). The steady technological improvements, together with the enhancement from better computer architectures, have contributed to a consistent performance improvement every year [2]. Fig. 1.2 depicts the comparisons of computer processor performance relative to a VAX-11/780 processor which are measured using the standard SPEC benchmarks over the years [2]. As shown in the figure, before 1980s, the growth in performance is largely driven by the technology advancement, which gives about 25% performance increase per year [2]. Then, this growth rate is improved to about 52% due to the introduction of more advanced architectural and computer organizational concepts, such as the emergence of the RISC (Reduced Instruction Set Computer) based microprocessors [2]. However, this trend begins to slow down again due to the constraints of power [3] [2]. As shown in Fig. 1.3, the power and power density on chip have been dramatically increased with the technology scaling down by the year of 2010. This is because as the technology process improves to a new generation, the increase in the number of transistors and the operating frequency overwhelms the decrease in load capacitance per transistor and the running voltage, which results in an overall growth in power density and energy [2]. For example, the Intel processor Core2 Duo consumes as high as 130W power which is about 20 times power of the 10 years ago Pentium processor in market [14]. The high power density causes severe reliability issues and inevitably makes the die cost unaffordable due to the high cooling requirements and costs. In order to avoid

reaching the power limit, the clock frequency growth with time has to be slowed down. This trend is shown in Fig. 1.4, where the clock frequency remains around $2GHz$ after year 2003 [2]. Instead of continuing the aggressive clock frequency scaling, the computer architects have proposed a new direction to achieve maximal performance under these tight constraints and budgets, *i.e.,* by employing more processor cores on chip for the computation tasks [4]. Consequently, the embedded systems have led to the multi-processor System-on-Chip (MPSoC) design and the high performance computer architectures have evolved into Chip Multiprocessor (CMP) platforms, which involve tens or hundreds processor elements, memory blocks, ASIC acceleration engines to be inter-connected together on chip. Each processing element performs its tasks in a parallel way taking the advantage of the parallelism either in task, thread or system level. As shown in Fig. 1.5, many recent chips have already switched to the paradigm of multi-core based platform for this purpose. For example, in [15], a 16-core heterogeneous digital baseband IC for MIMO 4G Software Defined Radio (SDR) is proposed. Their proposed NoC-based prototype doubles the throughput and consumes only 39% power over the previous MPSoC solutions. Another example is the Intel 80-tile Teraflops processor [16] which is a homogeneous NoC-based CMP platform and delivers up to 1.28 TFlops of performance. Recently, photonic on-chip network based multi-core systems have also been widely studied [17,18], which further attempts to optimize the traditional metal-based interconnect performance in terms of delay and power for future SoCs with thousands PEs.

1.2 NoCs for multi-core communication

For multicore based computing platform, an efficient way to manage the communication among the on-chip resources become critically important. The peer-to-peer interconnection consumes large wire area, which leads to large area and fan-outs [5]. The bus based architecture suffers from its limited bandwidth as well as scalability [5]. The scalability issue of these two schemes brings significant overhead in power and transmission delay when the chip feature size reduces beyond 45nm. To satisfy the communication requirements with hundreds or thousands processor elements (PEs), Network-on-Chip has been proposed as an efficient and scalable solution. Borrowing the concepts in Internet and wireless network, NoC

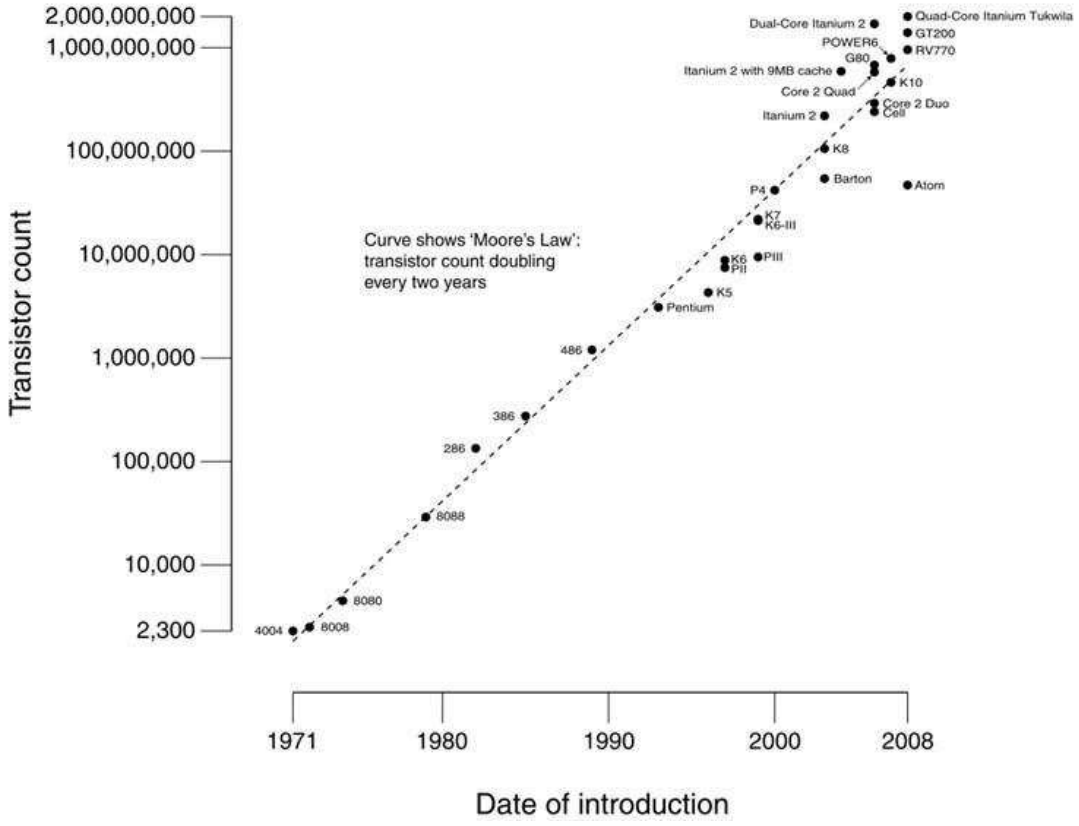

Figure 1.1: The plots of transistor counts against dates and Moore's Law (courtesy of [1])

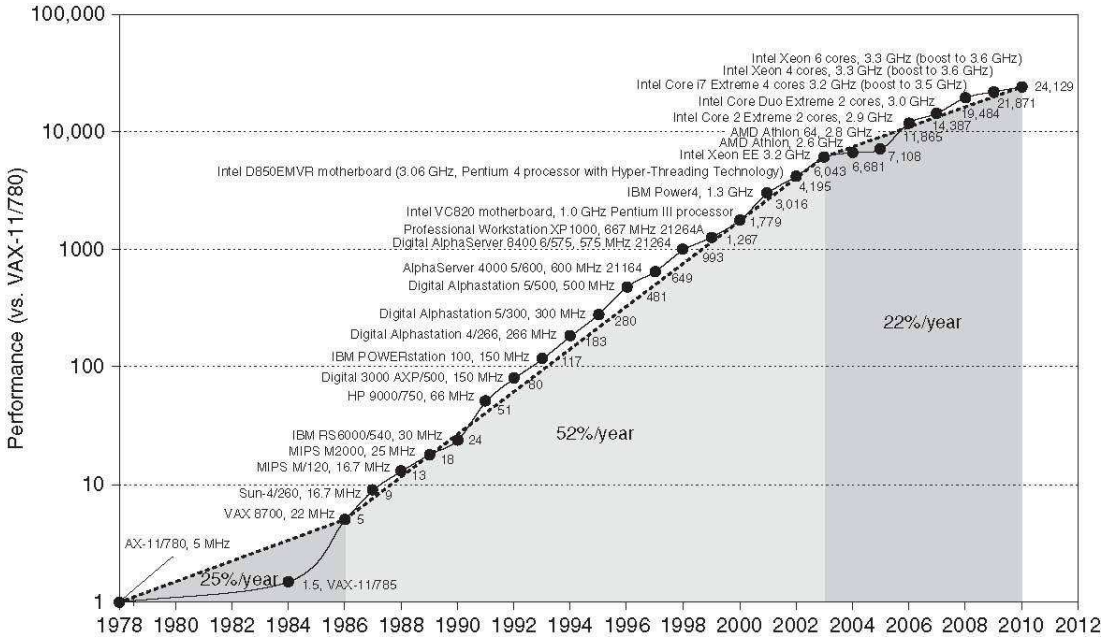

Figure 1.2: The illustration of the processor performance growth since the late 1970s (from [2])

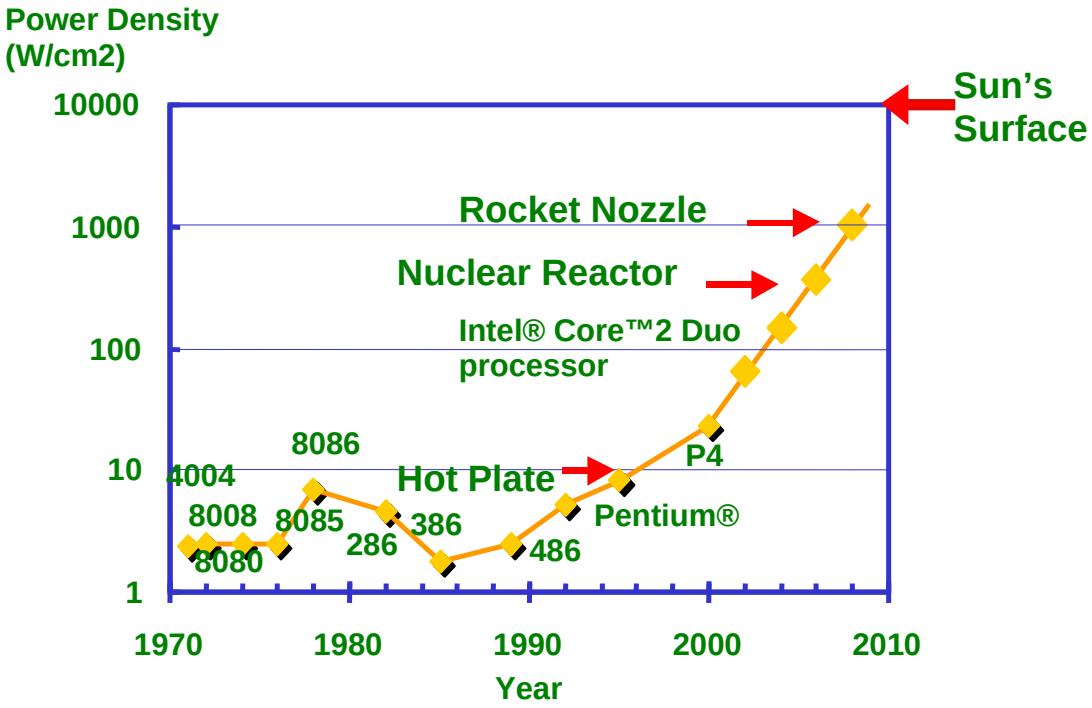

Figure 1.3: Power density trends in Intel's CPU (from [3])

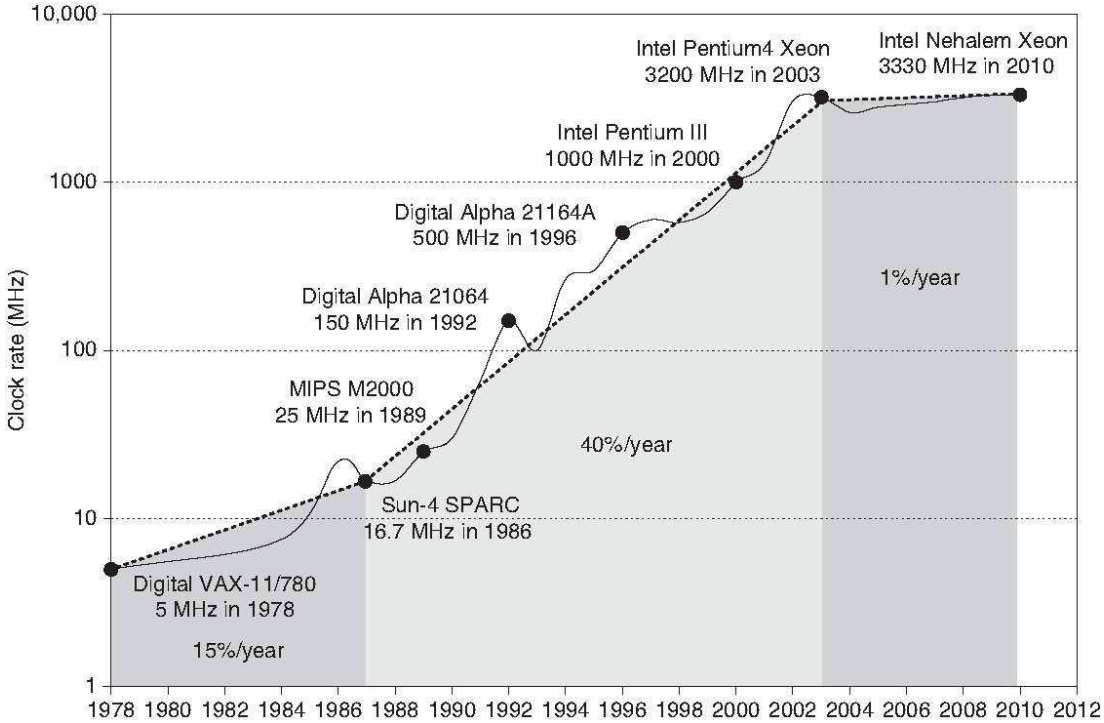

Figure 1.4: Growth in clock rate of microprocessors (from [2])

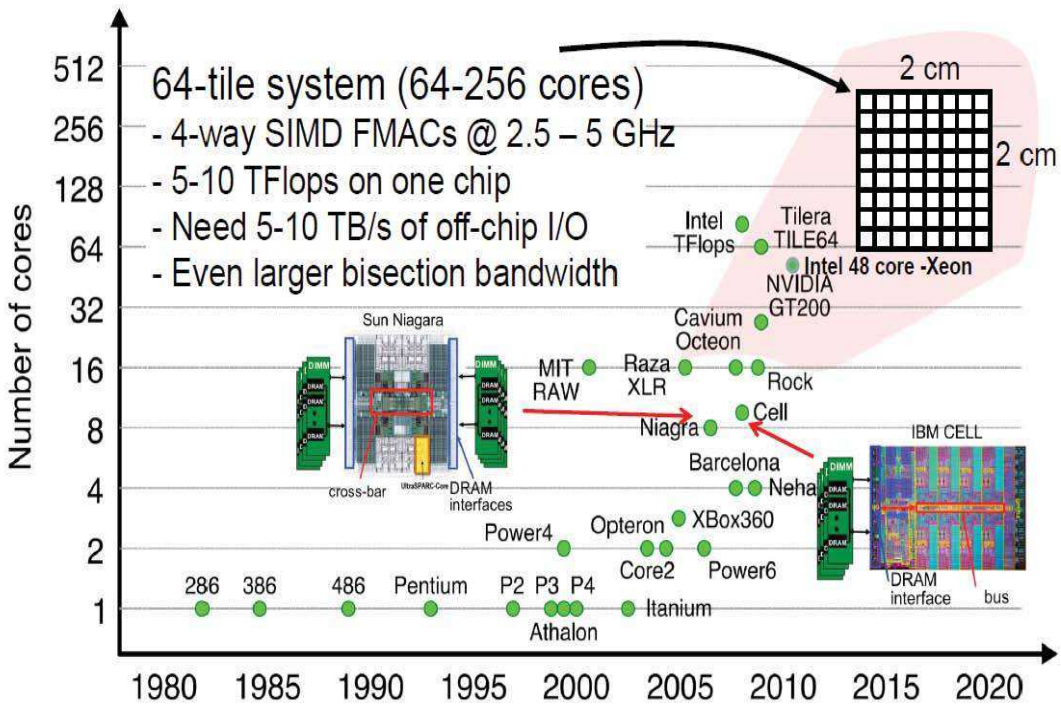

Figure 1.5: Manycore system roadmap for improving computing platform performance (from [4])

use routers to route packets instead of wires [6]. The latency and throughput performance is improved due to the higher bandwidth offered by the network. Meanwhile, the power consumption can be significantly reduced by breaking long links between the processors and avoiding high fan-outs in the outputs. In summary, Fig. 1.6 shows the trend of on-chip interconnection and compares the total wire length under different technology nodes [5]. As shown in the figure, for the technology nodes beyond 50nm, NoCs are more preferred over the other two paradigms in order to provide scalable communications for more than $1Km$ wire lengths.

In Fig. 1.7, we show an example of a MPSoC design using the mesh topology NoC for the video object plane decoder (VOPD) application [6]. The whole application is characterized by an application task graph [6], where the vertices in the graph represent certain computation tasks need to be performed and the edges indicate the communication bandwidth (MB/s) between two adjacent tasks. As shown in the figure, for the NoC-based VOPD platform, the whole system consists of twelve tiles organized in a rectangular mesh topology. Each tile is made up of a processor element (PE) and a router. The PE executes certain tasks in the application task graph while each router has five input/output ports that

are connected to the four neighboring routers as well as the local PE. At run time, the packets are routed based on the routing algorithm which is designed to determine the order of the routers to be traversed for a specific communication flow. For the NoC-based multicore system, besides the latency and throughput improvement, it also brings the following advantages:

1) High reliability: For Multicore systems, the complex system is highly susceptible to faults [19]. Compared to point-to-point dedicated links and buses, NoC can achieve higher reliability by providing redundant paths among the cores. If some of the routers fell into permanent or temporary faults, the other routers can be utilized to re-route the packets to the destinations and hence packet acceptance rate will not drop dramatically.

2) Modular design and IP re-use: NoC provides sufficient bandwidth for communication, while the processors can be designed without considering the network; therefore it supports modularity design and IP reuse [20]. Moreover, the global clock synchronization is not necessary in NoC which increase the overall system yield [14, 21].

3) Global asynchronous, locally synchronous (GALS) design: For multicore systems, it is difficult to distribute a single clock over thousands processor cores. To deal with these issues, NoC offers a good platform for the GALS design style [21] because each tile (processor elements and the router) can work separately within its own clock domain [14]. By employing GALS design, multiple Voltage-Frequency islands can be developed in different regions of NoC so as to achieve lower power consumption [14].

4) Power and area efficiency: Compared to the buses, the arbitration time for contention is much smaller as each router only needs to handle local contention scenario [22]. Therefore, large buffers to store the unserved packets are not needed in NoC routers, which result in a more compact router design and reduces the area/power overhead [22]. For power dissipation, because the buses are connected to all the PEs in the system, while the links in NoC only need to connect two neighboring routers (or a router and a PE) [22]. Therefore, with

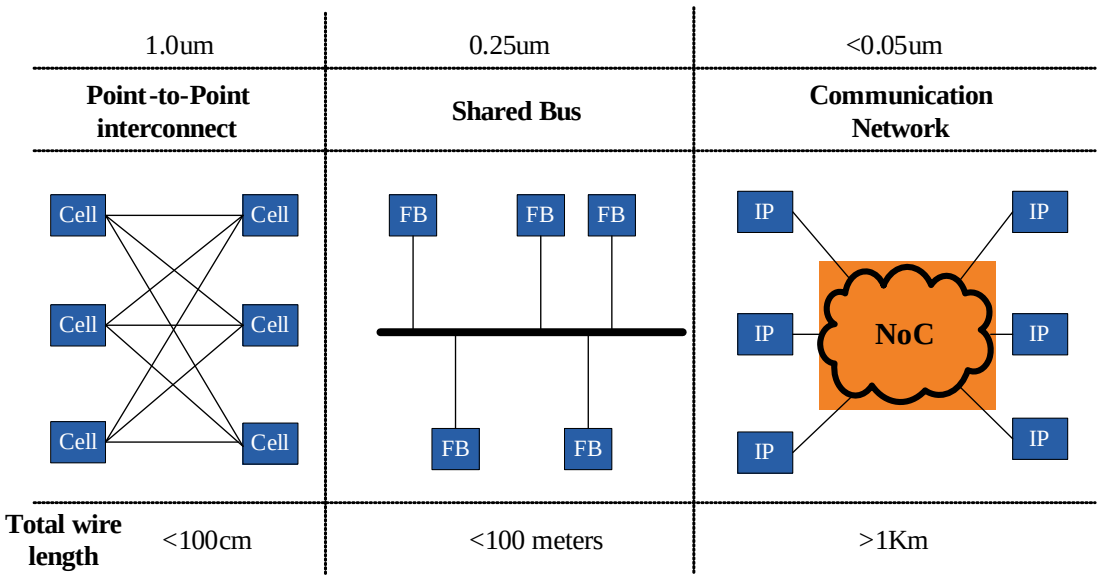

Figure 1.6: The trend of on-chip interconnections (from [5])

proper floorplanning, NoCs uses shorter wire length and occupies less load per transition [22]. Moreover, NoCs provide a variety of efficient power management strategies to further reduce power. This is because the NoC can be partitioned into sub-networks and each region can be powered-off or slowed down via dynamic voltage and frequency scaling (DVFS) individually [22]. High power efficiency can be achieved without significant degradation in the overall system performance.

1.3 Application-specific NoC design flow

In a typical NoC-based multicore system design, we begin with a specification of performance requirements combined with some cost constraints. These performance metrics, such as the latency/throughput, power consumption and hotspot temperature, drive the choice of NoC design parameters. A typical NoC synthesis flow works as follows (summarized from [6]):

1) Task scheduling and mapping: The first thing to determine is to allocate and schedule the tasks on the available processors. Usually a task graph is utilized to characterize the traffic patterns and the communication volumes of each traffic flow in the application (as shown in Fig. 1.7). Given the processors in the platform, task scheduling and mapping algorithms are developed to decide which processor that a specific task should be executed on as well as the order of the tasks to be executed on the same processor. In this step, bandwidth uti-

lization, total delay and power consumption are major design objectives while physical bandwidth as well as hard or soft deadline of some particular tasks are the constraints that need to be considered.

2) Core mapping: After the tasks are scheduled and mapped onto processors, the next step is to place these processors onto the NoC architecture. A core mapping algorithm is developed for this purpose. The core communication graph derived in the previous step determines the placement of tiles in this step. A mapping solution with high throughput, low latency and low power is usually desired while it should not exceed the capacities of the physical link bandwidth.

3) Routing algorithm design: After the task and processor mapping, routing algorithm is developed to decide the physical paths for sending the packets from the sources to the destinations. It will greatly affect the packet latency between the two cores as well as the overall chip power and thermal profile. For the routing algorithm design, one important issue is to avoid deadlock. The deadlock refers to the situation that the whole system stalls due to the circular dependencies [23]. More specifically, for the deadlock scenario, it happens at run time, where flits from some packets occupies some resources in the router (such as the buffer). At the same time, they request to use other resources (such as the buffer in the downstream node), so the dependencies of the channels may have chances to form a cycle. In this case, all these packets are stalled in place and can not proceed to the destination anymore [23].

As there are a lot of possible design choices in each of the three steps above, NoC-based multicore system design produces a large space to be explored. Therefore, it is of utmost importance to provide an accurate performance evaluation with respect to the specific configurations in the synthesis inner loop. Both analytical models and simulations can be used in the NoC performance evaluations. To fully understand and model the details of the network situations occurred at run time, NoC simulators are developed and widely adopted with high fidelity. On the other hand, since NoC designs have many power, area and latency trade-offs in topology, task and core mapping algorithms *etc.*, analytic models have also been deployed to allow fast design space explorations [24]. In general, it is more reasonable to work with simple analytical models first in the synthesis loops, while

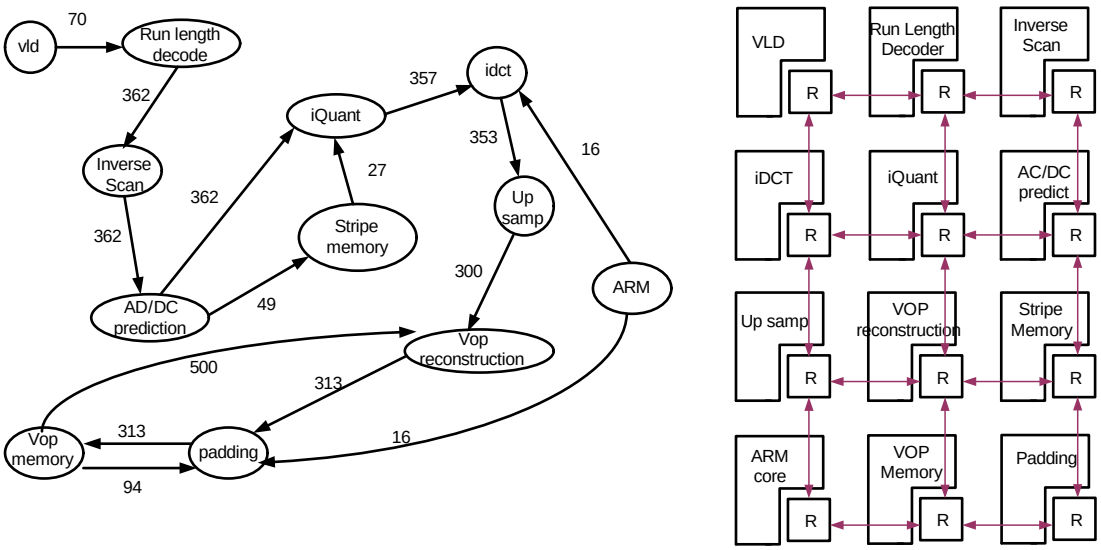

Figure 1.7: The NoC-based Video Object Plane Decoder (VOPD) platform in [6]

more detailed simulations become necessary to accurately characterize the exact performance of the network after only a few candidates being remained [24]. In Fig. 1.8, we summarized the synthesis flow for the NoC design. The usage of analytical models is highlighted within the inner loop in the figure.

1.4 NoC characterization

In order to better understand the terminologies, concepts and algorithms developed in this thesis, in the following, we briefly review the basic ideas and the models that characterizes an NoC platform.

1.4.1 NoC topology

Network topology refers to the arrangement of various elements (links, nodes, *etc.*) of the network [20]. Essentially, it is the topological structure of the multi-core platform and is dependent on the placement of the network's components, including the locations of the processors and the routers [20]. There are various NoC topologies, such as mesh, torus [25], butterfly [26], 3D-mesh [27] and fat trees [28].

In this thesis, we assumed the underlying NoC system is composed of a 2D mesh network (as shown in Fig. 1.7). The reason for using the 2D mesh network is due to its regularity and layout efficiency on silicon surface [25]. Moreover, the

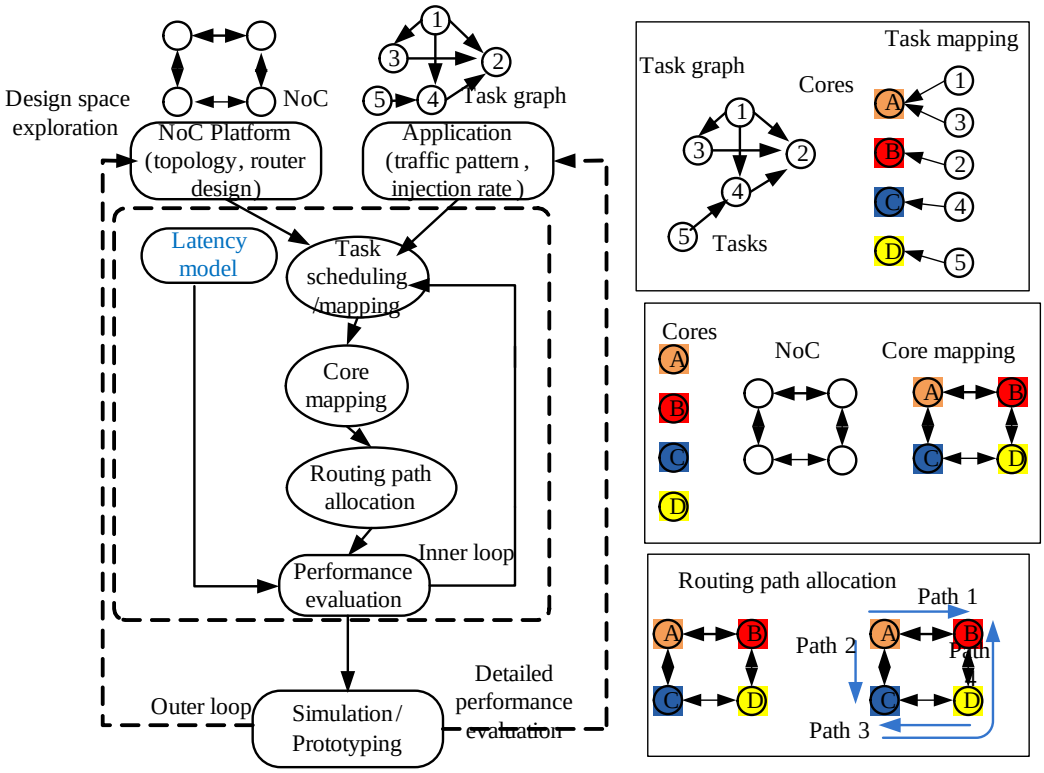

Figure 1.8: NoC synthesis flow for application specific multicore systems

2D mesh topology also matches well with the current IC manufacturing technology for the layout consideration, especifally for most IC components which have rectangular shape [14]. Therefore, this topology has attracted wide attention in most state-of-the-art NoC-based multicore prototypes (*e.g.,* MITs 16-tile RAW chip [29] and Intels 80-tile TFLOPS chip [16]).

Another advantage of the mesh based topology is the high scalability to merge or combine building blocks which are developed with regular shapes [14]. When the complexity of the embedded systems is increasing, more PEs are trying to be put together on the chip. With regular shape, the additional PE blocks can be easily integrated on the original design [14] which eases the voltage/frequency island based control on NoC.

1.4.2 NoC switching technique

In general, based on the flow control granularity, the NoC routers can be classified into three types, namely the circuit switching, virtual cut-through switching and the wormhole switching [6]. The reviews are done based on [6, 14, 30]:

In the circuit switching paradigm, two PEs set up a specific communications

channel (named as circuit path) in NoC first before they begin to transfer packets to each other. The circuit switching ensures the bandwidth for the channel settled and keeps connected during the whole communication period of the specific flow [30]. However, it is sometimes inefficient in using the channel bandwidth because the unused links reserved for one connection cannot be used by others when the circuit is set up [6, 30].

In virtual cut-through switching [31], the buffers are designed to be capable of storing the whole maximum packet. However, the whole packet is only stored into a router buffer if the downstream router buffer is already occupied by other packets [14]. Otherwise, the flits once arrived at the current buffer can be routed directly without the need to wait for the arriving of other flits in the same packet [14]. Hence, in virtual cut-through switching, if the packet stall [20] happens due to the failure of allocating a downsteam channel, the packet stays in the current node will not block any other packets [14]. Compared to the circuit switching, as the flits can be forwarded immediately, the network latency under no congestion is reduced; however, the virtual cut-through switching still requires the buffer size to be large enough in order to store the whole packet under congestion [6, 14, 30].

In order to overcome the limitations in the circuit and virtual cut through switching, the wormhole switching techniques have been proposed and widely used in the communication networks [30] as it requires fewer buffer resources than previous two techniques [6, 30]. In particular, in wormhole NoC, each message consists of several packets. Furthermore, the packet is divided into several flits, which are the minimal flow control units in the routing. The header flit is utilized to settle the routing paths in the routers, while the body and tail flits simply follow the paths reserved by its header. When the tail flit leaves the router, it will release all the resources it reserved for the packet so that the consequent packets can use them again [6, 30]. Therefore, one major advantage of wormhole routing is that it does not need a large enough buffer to hold the whole packet, which drastically reduces the overall latency [14].

1.4.3 NoC router design

The typical structure of an on-chip router for a mesh NoC is shown in Fig. 1.9, where we show the basic control and data path with four virtual channels (VCs)

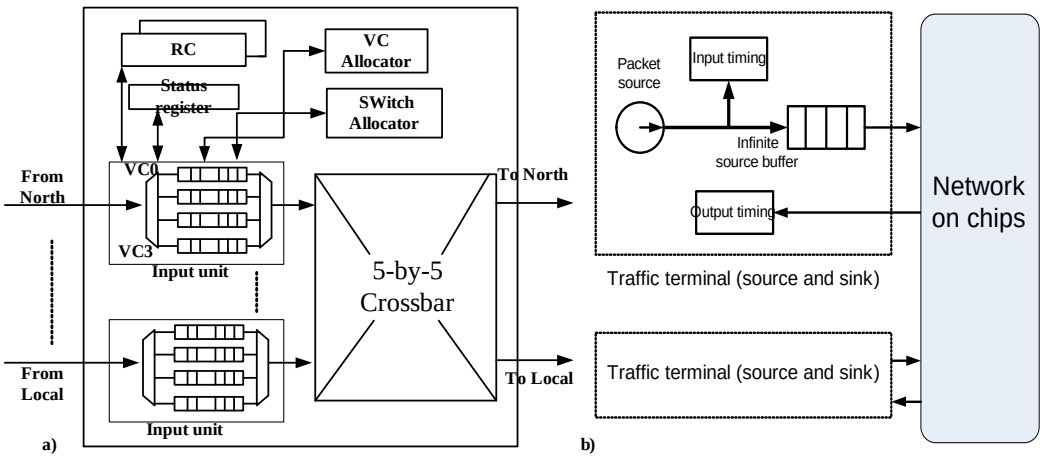

Figure 1.9: The NoC router architecture with virtual channels and the open-loop source queuing delay measurement

in each input port [6, 20]. For the wormhole router, it can be viewed as a special case of virtual channel routers where the number of VCs per input port equals to one. As shown in Fig. 1.9, the router control path usually consists of the routing computation (RC) module, the virtual channel allocation (VA) module and the switch allocation (SA) module. Its data path usually consists of the buffers (virtual channels or a single buffer), the switching crossbar and the output registers. Of note, the RC module is used to compute the output port according to the destination address recorded in the header flit in front of the buffer. The VA module is used to allocate the downstream virtual channels to the packet in the current virtual channel buffer. After a virtual channel is allocated to the packet, the switch allocation module works to arbitrate for the usage of the switch fabric (crossbar) entries among the packets. On top of this baseline architecture, many modifications on the router datapath and control path have been proposed to reduce the power consumption and latency (*e.g.,* the speculative router [20] and the lookahead router with bypass architectures [32]).

The network interface (NI) is needed between the router and PE to convert and transfer messages. The standard measurement setup for interconnection networks is shown in Fig. 1.9-b [20]. To measure the performance of an interconnection network, we need to attach terminals (*i.e.,* PEs) to the local port of the adjacent router in the network. The NI is usually modeled with infinite buffer size to isolate the NoCs from the processors [20]. It is important that in this open-loop measurement, the monitors in NI is placed ahead of the source queue instead of after the queue (the monitor records the injection and ejection times of the

packets to PE) [20]. In this way, the packets that have been generated and are still waiting to be injected into the network are considered [20]. Therefore, the overall packet latency under this set-up will not only include the NoC traverse time but also the source queuing time [20]. Various workload can be applied to the NI, including the traffics generated from real processor models, the trace-based workload as well as the synthetic workload without using any processor information [20].

1.5 Thesis overview

In our work, we investigated the issues for a high performance NoC design, from analytical performance modeling to the routing algorithm design and NoC architecture optimization. In particular, we looked at the following areas: NoC performance modeling for the design space exploration (Chapter 2); routing algorithm design for the thermal-awareness and reliability objectives (Chapter 3 and 4); flexible NoC architecture design using self-reconfigurable bi-directional channels (Chapter 5 and Chapter 6). We aim to address several key problems in NoC design from both the algorithmic point of view as well as the hardware and architectural level optimization. To be more precise, the outline and contributions of this thesis are summarized below:

Chapter 2: SVR-NoC: A learning based NoC latency model

In this chapter, instead of using conventional queuing-theory-based NoC latency model, which have several assumptions that comprise the overall prediction accuracy. We proposed a learning based NoC latency regression model, namely SVR-NoC, to accurately evaluate a candidate design in the inner loop for the design space exploration. Compared to the previous models, we showed that better accuracy over the queuing models and at the same time $70X - 102X$ speedup over the detailed simulations can be achieved using the SVR-NoC model on both the synthetic traffic patterns and real application traces. Therefore, the SVR-NoC can benefit the exploration of numerous design configurations in the offline phase.

Chapter 3: A thermal-aware routing algorithm for application-specific Network-on-Chips (NoCs)

For NoC-based multi-core systems, the routing algorithm significantly affects the overall performance and needs to be tackled in the offline design phase to meet

the certain design constraints. Among all the routing considerations, temperature is one of the most critical ones as the uneven temperature across the chip will introduce thermal hotspots and degrade the performance dramatically. Towards this end, in this chapter, we propose an offline thermal-aware routing algorithm to evenly distributed the traffic across the chip so to reduce the hotspot temperature. Specifically, we propose a deadlock free adaptive routing algorithm which provides maximal number of paths to route packets for the given application. Then, a linear programming based algorithm is used to find the optimal ratio to send packets among the paths. We show that as much as $10\% - 20\%$ peak energy reduction can be achieved over a set of applications while the latency/throughput performance is maintained by using the proposed method.

Chapter 4: Fault-tolerant NoC routing algorithms design

In this chapter, we investigate and propose a highly resilient routing algorithm to tackle the router and link faults at run time. More specifically, we classify the permanent faults in NoC into two types (*i.e.,* link faults and buffer faults). For the link faults, a highly resilient routing algorithm is used to re-route the packets from faulty links. While for the buffer faults, we propose two new schemes, namely dynamic buffer swapping and dynamic MUX swapping to handle the errors in the buffers and crossbar Muxes, respectively. We show that, higher packet acceptance rate as well as better latency and throughput performance can be achieved for a set of test traffics.

Chapter 5: FSNoC: A flit level speedup scheme for NoCs using self-reconfigurable bi-directional channels

Besides the optimization in the algorithm level, we also explored to improve NoC performance from the architectural level. In this chapter, we propose FSNoC, a new NoC router architecture that supports switching two flits from the same packet simultaneously by using the bi-directional channels. Compared to previous router architectures using bi-directional links, better link bandwidth utilization can be achieved and therefore FSNoC can lead to higher performance in latency and throughput. The channel direction control protocol as well as the router micro-architecture which supports flit-level parallel transmission have been proposed. We demonstrate the performance improvement of FSNoC using both synthetic traffic patterns as well as the traces from realistic applications. The hardware overhead of FSNoC in terms of area and power is also reported and

analyzed in detail in this chapter.

Chapter 6: A traffic-aware adaptive routing algorithm on a highly flexible NoC architecture

In this chapter, we aims to add more flexible in the overall NoC architecture and propose a new platform which consists of i) self-reconfigurable bi-directional channels, ii) express virtual channels and iii) regional hub routers to improve the system performance. A fitness-based and traffic-aware adaptive routing algorithm is designed which is suitable for the proposed platform and chooses the routing path dynamically to adapt the traffic conditions at run time. Combining the routing algorithm and the platform, more than 80% improvement in saturation throughput can be obtained, while involving less than 15% overhead in power dissipation.

Chapter 7: Conclusion and future work

This chapter summarizes the works done in the whole thesis and discusses several future research directions.

CHAPTER 2

SVR-NOC: A LEARNING BASED NOC LATENCY PERFORMANCE MODEL

In this Chapter, we propose SVR-NoC, a learning based Network-on-Chip (NoC) latency model using support vector regression (SVR). Different from the state-of-the-art NoC analytical models, which use queuing models to compute the average channel waiting time and the source queuing time, the proposed SVR-NoC model predicts the NoC latency based on learning the typical training data. More specifically, given the application communication graph, the NoC architecture and the routing algorithm, we first analyze the links dependency and then determines the ordering of latency analysis. The channel and source queue waiting times are then estimated using a new generalized $GE/G/1/K$ queuing model, which can tackle bursty arrival times with general service time distributions. To improve the prediction accuracy, the queuing theory based delays are included as one of the features in the learning process. We propose a systematic learning framework that uses the kernel-based support vector regression method to collect training data and predict the traffic flow latency. The proposed learning-based model can be used to analyze various traffic scenarios for NoC platforms with arbitrary buffer and packet length. Experimental results on both synthetic and real applications demonstrate the accuracy and scalability of the proposed SVR-NoC model as well as a $120X$ speedup over simulation-based evaluation methods.

2.1 Introduction

The NoC complexity as well as its tight requirements on the power, latency, and throughput have become major challenges in the design of NoC-based multi-core systems [35]. As shown in Fig. 2.1, a typical NoC-based system design requires many synthesis steps including task allocation, mapping, core placement and routing. Specifically, based on the pre-characterized application traffic, the designers first need to schedule and map the tasks on the available processor elements (PEs). After the task scheduling and mapping are done, core placement

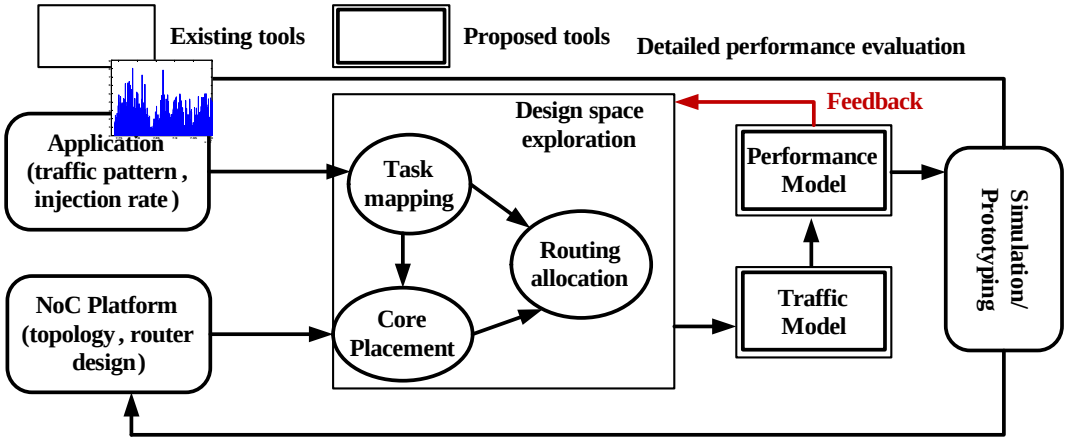

Figure 2.1: The NoC performance model used in the synthesis inner loop for large design space exploration

Table 2.1: Summary and comparison of NoC latency models

	Queuing theory based analytical models						**Learning based**
Models	[34]	[33, 35]	[24]	[36]	[7]	Proposed	SVR-NoC
Queue	M/M/1	M/G/1/K	G/G/1/∞	M/M/m/K	GE/G/1/∞	GE/G/1/K	N/A
	Application related models						
Arrival	Poisson	Poisson	General	Poisson	GE	GE	General
Service	Memoryless	General	General	Memoryless	General	General	General
Traffic pattern	Arbitrary	Arbitrary	Arbitrary	Arbitrary	Uniform	Arbitrary	Arbitrary
	NoC architecture related models						
Buffer size	1 flit	K packets	B flits	1 flit	1 flit	Arbitrary	Arbitrary
PB ratio[1]	m ($m \gg 1$)	< 1	Arbitrary	m ($m \gg 1$)	m ($m \gg 1$)	Arbitrary	Arbitrary
Arbitration	Round-robin	Round-robin	Round-robin	Fixed priority	Round-robin	Round-robin	Round-robin

and routing path allocation need to be explored. Each of the above steps can produce numerous design choices. Therefore, a performance analysis tool is needed to evaluate whether the chosen NoC configuration for the input application leads a better design over others while satisfying the design constraints at the same time. Detailed network simulations can provide performance evaluation results with high fidelity. However, they suffer from long evaluation times and so are only suitable for estimating a small subset of alternatives in the final prototyping stage. Because of this, NoC performance models are widely adopted to guide the pruning of the design space during the system synthesis [33].

Among all the NoC performance metrics, latency is recognized as one of the

most critical design parameters since it determines the whole system throughput under specific workloads [24]. In this work, we propose a latency model to predict the average delay of flows in an NoC-based multi-core system for the design space exploration. In order to derive a latency model, most previous researches are based on the queuing-theory formalisms and treat each input channel in the NoC router as an $M/M/1$ [34], $M/G/1/N$ [33], $G/G/1$ [24]. Indeed, these models provide accurate performance estimations when the following assumptions hold: i) The packet length satisfies an exponential distribution, and therefore the packet service time in the router is exponentially distributed as well [34]. ii) The traffic inter-arrival time is assumed to follow a Poisson distribution at all traffic sources [33, 35]. However, it has been observed that in many NoC systems, the behavior of the traffic follows the fractal/long-range-dependent (LRD) pattern [37] and the distributions of the service time are correlated as well. Consequently, the accuracy of the queuing theory-based model is compromised in these cases.

In this chapter, we attempt to develop an NoC latency model which is suitable for the synthesis inner loop and has higher prediction accuracy by using a new approach based on the machine learning techniques. More specifically, we first propose a new queuing-theory-based delay evaluation methodology which can work for a variety of NoC configurations and traffic scenarios. The proposed performance queuing (PQ) model is based on a $G/G/1/K$ queuing formalism and generalizes the previous NoC PQ models as follows: i) The existing traffic arrival modeling using Poisson approximations is extended to a generalized exponential (GE) packet inter-arrival distribution which can account for burst traffic patterns. ii) The packet service process within each router is modeled with a general distribution to account for the service time correlation between routers and traffic flows. iii) A more general NoC architecture model is used so that routers with finite buffer depth are accommodated, thus enabling the consideration of arbitrary buffer depth and packet length combinations. iv) By considering the link dependencies, the proposed framework is completely generic and can be applied to any NoC topology with different task mapping and routing algorithms. Then, to relax some of the assumptions in the PQ model (such as the GE traffic arrival process) and further improve the modeling accuracy, we propose SVR-NoC, which is a support vector regression (SVR) based NoC latency model. In SVR-NoC, the delay predictions based on enhanced queuing-theory-based PQ model are included

as part of the features in the learning process. In the training stage, the training data-set is formed by collecting the latency simulation results of the same NoC platform on various synthetic traffic patterns. We employ *Support vector regression* [38] techniques to learn the channel queuing and the source queuing models as functions of their feature sets, respectively. During the learning process, *cross validation* is used to avoid training data over-fitting. In the prediction stage, the learned SVR model is used to estimate the average waiting time at each input buffer channel and then the overall packet flow latency for the new application patterns.

The rest of the chapter is organized as follows. In Section 2.2, we review the previous arts and highlight our contributions. In Section 2.3, we present the proposed generic queuing-theory based latency model. Section 2.4 details the learning-based SVR-NoC latency model. The experimental results of the proposed NoC latency model on both synthetic and real applications are shown in section 2.5. Finally, Section 2.6 concludes this chapter.

2.2 Background

The analytical models for evaluating the NoC average latency can be classified into three groups: probabilistic models [39], network calculus models [40], and queuing theory models [24,33,35]. In [39], a probabilistic analysis framework was developed to model a single wormhole router performance. However, additional effort is needed to extend to network of routers. In [40], the network calculus approach was adopted to characterize the NoC performance. However, the average delay prediction error is larger than that of the queuing models [40]. Therefore, most of the previous efforts are based on queuing models to evaluate the NoC delay.

For the general class of queuing-theory-based NoC models, most of the early works consider the modeling of wormhole (WH) routers under the assumption of Poisson arrival time distribution and memoryless packet service time distribution. For example, in [34], an M/M/1 approximation of link delay is used to analyze the capacity and flow allocation. Although generally tractable, the accuracy of M/M/1 models can be significantly compromised as the assumption of exponential arrival and service time distributions may not hold in many real

Table 2.2: Parameters and notations in NoC latency model

Parameters	Description
H_s	Service time for head flit (including the link transfer)
m	Average packet size (flits)
B	Buffer size in each channel (flits)
$f_{s,d}$	Communication flow from source s to destination d
d_f	Length (number of hops) of flow f
l_{ab}	Link channel connecting router a and b
$P_f\,(P_{s,d})$	Set of links that form the routing path of flow $f_{s,d}$
$F_{l_{ab}}$	Aggregate set of flows sharing link l_{ab}
l_i^f	Link that resides in the i^{th} hop of flow f
C_f^2	SCV of the packet inter- arrival time of f
$\lambda_{l_{ab}}$	Mean packet arrival rate at link l_{ab}
$q_{l_i^f}$	Delay for a flit to reach the head of buffer in link l_i^f
$h_{l_i^f}$	Delay for a packet head to acquire the link l_i^f
$\eta_{l_i^f}$	Header flit transfer time over link l_i^f
$s_{l_i^f}/s_{l_{ab}}$	Service time for a packet that travels link l_i^f/ l_{ab}
$z_{l_i^f}$	Time that a header flit reaches the point where the accumulated buffer space can hold the whole packet
v_s	Source queuing time at time s
$L_f\,(L_{s,d})$	Average latency of flow $f_{s,d}$ (cycles)

applications [37, 41]. Several works have been proposed to improve the estimation accuracy by generalizing the arrival and service time distributions. In [42], an analytical model based on M/G/1/K queue is proposed to account for finite size input buffers in local area networks (LANs). However, this model is based on the Laplace-Stieltjes transform and is too complicated to be used in the NoC synthesis loop. In [35], an M/G/1 based latency model for NoC analysis is proposed. It only assumes that the arrival rate of the header flits (as opposed to the entire packet) follows a Poisson distribution. In [24], a fixed-priority G/G/1 based NoC latency model, which attempts to model the burst arrival times with a 2-state Markov-modulated Poisson process (MMPP), was proposed. However, this approach targets a specific priority-based router architecture, while many NoC routers may utilize a more fair arbitration such as the round robin (RR) scheme. In [36], an M/M/m/K queue-based analytical model is proposed to analyze the delay of NoCs with variable virtual channels per link. This approach assumes negligible flit buffers (*i.e.,* single flit buffer) such that a packet reaches its destination before its tail leaves the source host. In [33], an $M/G/1/N$ queuing mode is proposed for both wormhole and virtual channel NoCs. However, this

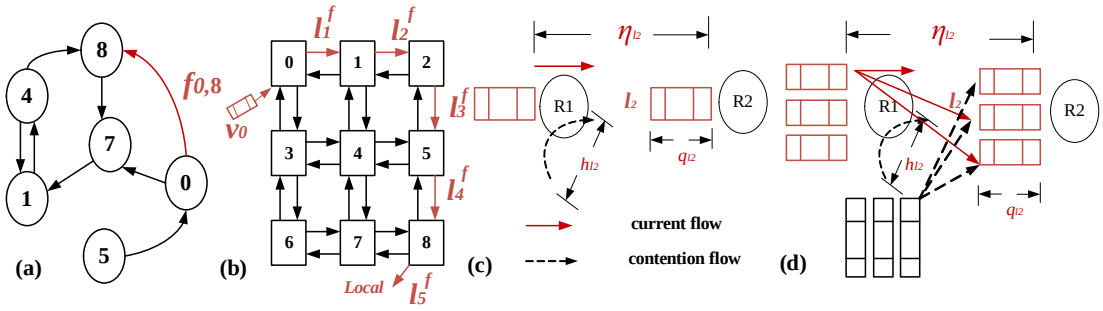

Figure 2.2: An example showing the flow delay: a) the core communication graph, b) the links involved to calculate the end-to-end delay of flow $f_{0,8}$ c) the acquisition time h and transfer time η for link $l_{1,2}$ in WH model, d) h and η in VC model

approach assumes that the granularity of the buffers is in terms of packets instead of flits and therefore a single channel buffer can hold up to N packets during the analysis. This may not be the case for NoCs whose buffers are rather small (only several flits) to save area and power [24].

Machine learning is a technique that has been extensively used in the domain of pattern recognition or artificial intelligence where it is usually difficult to derive the exact mathematical relationship between the outputs and inputs in these problem formulations [38]. In NoC performance modeling and analysis, most of the learning models focus on using learning-based model to improve the area/power model accuracy. For example, in [43], the multivariate adaptive regression splines (MARS) technique is used to develop a non-parametric NoC router power and area regression model. Compared to the conventional ORION2.0 model [44], the learning-based model improves the prediction accuracy over a variety of NoC implementations. Later, in [45], the model accuracy is further enhanced by explicitly modeling of control and data paths in the regression analysis. Besides area/power modeling, the learning techniques have also been used in optimizing NoC runtime performance, such as the reinforcement learning based DVFS control [46], the Q-learning based congestion-aware routing [47] and the neural network based optimal dynamic routing [48]. However, for the latency performance metric analysis, machine learning techniques for improving modeling accuracy over the conventional queuing model have not been thoroughly studied yet.

In this chapter, we propose a new learning-based NoC latency model which generalizes the previous work by considering: i) the arrival traffic burstiness, ii)

the general service time distribution, iii) the finite buffer depth and arbitrary packet length. For clarity purposes, in Table 2.1, we summarize and compare our proposed queuing model and the learning-based SVR-NoC model against other models proposed to date [2]. As shown in the table, our proposed model offers a much broader coverage for various temporal and spatial traffic patterns, as well as NoC architectures. This provides more flexibility for the designers to explore the NoC design space.

To the best of our knowledge, this chapter brings the following new contributions over the previous efforts:

- We proposed a new queuing-theory based NoC traffic model which is topology-independent and can be used to analyze a variety of traffic scenarios as well as arbitrary buffer size and packet length combinations.
- In addition to the proposed queuing model, we propose and develop a learning-based framework for NoC latency analysis. The model has fewer assumptions related to the packet length and traffic distribution as well as the router architectures.
- We show the accuracy and scalability of the proposed SVR model using both the synthetic traffic and real applications. The speedup of the learning model over the conventional simulations can significantly benefit the NoC synthesis and optimization process.

2.3 Proposed NoC analytical model

2.3.1 Basic assumptions and notations

We assume that the target applications have been scheduled and mapped onto the target NoC platform and the source and destination tile addresses for each specific flow f in the application are known. Also, borrowing from the idea of modeling bursty traffic in hyper-cube multi-computers [49], we assume that the packet inter-arrival times of flow f have been characterized using a general exponential (GE) distribution [7] (discussed later in Section 2.3.4) with mean λ_f^{-1} and a square

[2] In the table, PB ratio is defined as the ratio of average packet size (m flits) to the buffer depth (B flits)

coefficient of variation (SCV) C_f^2. Therefore, the (λ_f^{-1}, C_f^2) characterization of the traffic model is an input to our analytical framework. Moreover, a deadlock-free and deterministic routing algorithm is used to guarantee that no cycles are formed by the link dependencies. Without loss of generality, in this work we use X-Y routing. Other deadlock-free and deterministic routing schemes can also be used. We also adopt a wormhole router architecture, where there exists a single buffer at each input port. For simplicity, we assume that the packets have a fixed size of m (flits) as in [41]. However, this assumption can be relaxed to cover arbitrary packet length distribution. Other assumptions are that the traffic sources (i.e., the source PEs) have an infinite queue size and the destinations immediately consume the arriving flits. To facilitate the discussion, the symbols in Table 2.2 are used consistently in this chapter, which follows the name conventions in [36, 42].

2.3.2 End-to-end delay formulation

In a WH or VC NoC, the end-to-end flow latency $L_{s,d}$ of a specific flow $f_{s,d}$ (shown in Fig. 2.2-a and -b) is made up of three parts [36, 42]: 1) the queuing time at the source s (*i.e.*, v_s), 2) the packet transfer time (*i.e.*, $\eta_{s,d}$) and 3) the path acquisition time (*i.e.*, $h_{s,d}$). It is expressed as [36]:

$$L_{s,d} = v_s + \eta_{s,d} + h_{s,d} \tag{2.1}$$

In order to calculate $h_{s,d}$, we need to consider the path acquisition time of every link l_i^f residing in the routing path of f (Fig 2.2-b shows the links used for the flow $f_{0,8}$) [36] and therefore: $h_{s,d} = \sum_{i=1}^{d_f} h_{l_i^f}$, where $h_{l_i^f}$ is the time for a packet header to contend a channel (in WH routing) or a VC (in VC routing) with other flows for link l_i^f and d_f is routing path length (number of hops) of f.

If we denote the time to transmit the header flit by $\eta_{l_i^f}$, then the packet transfer time $\eta_{s,d}$ can be rewritten as [36]:

$$\eta_{s,d} = \sum_{i=1}^{d_f} \eta_{l_i^f} + (m-1) \tag{2.2}$$

where the first term denotes the header flit transmission time over the network, and the second term approximates the packet serialization time of the body and

tail flits. The notations of the queuing delay v_s, $h_{l_i^f}$and $\eta_{l_i^f}$ are also illustrated in Fig.2.2 for both the WH and VC NoCs.

In order to derive $L_{s,d}$, $\eta_{l_i^f}$ and $h_{l_i^f}$ for each link, which are the elements of the routing path of $f_{s,d}$, need to be obtained. Two issues hereby arise. First, it is important to determine the order of analyzing the links of the path. Second, the detail model to obtain η and h of each link has to be developed. In particular, we need to differentiate between the VC routing and the WH routing, which are shown in Fig. 2.2-c and -d, respectively. As shown in the figure, the VC router complicates the η and h modeling in the following two aspects: i) For the link contention time h, the packet can randomly choose among V VCs instead of 1 and therefore the simple single server queuing model in WH models can not be used directly. ii) For the flit transfer time η, we need to consider the additional time required due to the flow multiplexing among the VCs over the same physical link as described in [36].

In the following subsections, we first present the procedure to determine the orders of the analysis of the links. After that, we will then elaborate on the formulation of the WH and VC router models.

2.3.3 Link dependency analysis

To account for the impact of the congestion at the downstream routers on the blocking of a particular router, it is important to obtain the dependency among all the links. Fig. 2.2-b shows an example. Due to back-pressure, the waiting times η_2^f and h_2^f of link l_{12} (*i.e.,* l_2^f) will affect the time to serve a packet in the buffer head of link l_{01} (*i.e.,* l_1^f). Therefore, l_{01} is said to be dependent on l_{12}. Similarly, η_2^f and h_2^f depend on those downstream links of l_{12} which is determined by the application mapping and routing. Thus it is required to finish the analysis of all the downstream links of link l_{01} in the flow f before we can calculate the queuing service and waiting time of link l_{01}. To obtain the link dependencies, a link dependency graph (LDG) is built first and the topological sort algorithm as used in [42] is then used to order the links. The detail of link dependency analysis is presented in Algorithm 2. The vertices in LDG correspond to the link channels in the NoC while a directed edge joining two vertices reflects that there is a dependency between these two links *(e.g., l_{01} and l_{12} in* Fig. 2.2*).* When the

routing or the task mapping solution changes, the LDG needs to be rebuilt. The LDG is built by checking every flow f in the application communication set F. An edge is added between two vertices of the LDG if there exists a flow $f_{s,d} \in F$ that the two links corresponding to the two vertices are two adjacent links on the routing path. The order of the link analysis are then obtained by applying the topological sort algorithm on the LDG.

Algorithm 1 Link dependency analysis

```
1: Input: F the application flow set
2: Output: G the ordered list of links for queuing analysis
3: Container: LDG = (V, E) the link dependency graph
4: for all link l_ab do
5:    LDG.addnote(l_ab)
6: end for{initialize the LDG}
7: for all flow f ∈ F do
8:    P_f = routing_function(f);
9:    d_f = length(P_f);
10:   for i = 1 : d_f − 1 do
11:      l_up = P_f(i); {the upstream link}
12:      l_down = P_f(i + 1); {the downstream link}
13:      LDG.addedge(l_down, l_up);
14:   end for{considering all the links in the path of f}
15: end for{considering all flows in application}
16: return G = topological_sort(LDG);
```

2.3.4 GE-type traffic modeling

Many applications in NoCs show burst patterns of traffic over a wide range of time scales. Therefore, the Generalized Exponential (GE) distribution is utilized to model the arrival traffic at the source PEs and the links [7]. In the following, we briefly review the GE type distribution proposed in [7,49]. Under GE distribution, the cumulative distribution function (cdf) of the inter-arrival time X is given by [7]:

$$F(t) = P(X \leq t) = 1 - \tau e^{-\tau\lambda t},\ t \geq 0 \tag{2.3}$$

where $\tau = 2/(1 + C^2)$ and (λ^{-1}, C^2) are the mean and square coefficient of variation (SCV) of X. The GE packet generation process is shown in Fig. 2.3 from [7, 49]. As shown in the figure, a packet experiences a zero service time to reach the departure point (*i.e.,* the point to generate the packet) with probability $1-\tau$. For the rest (*i.e.,* with probability τ), the packet needs to traverse the system

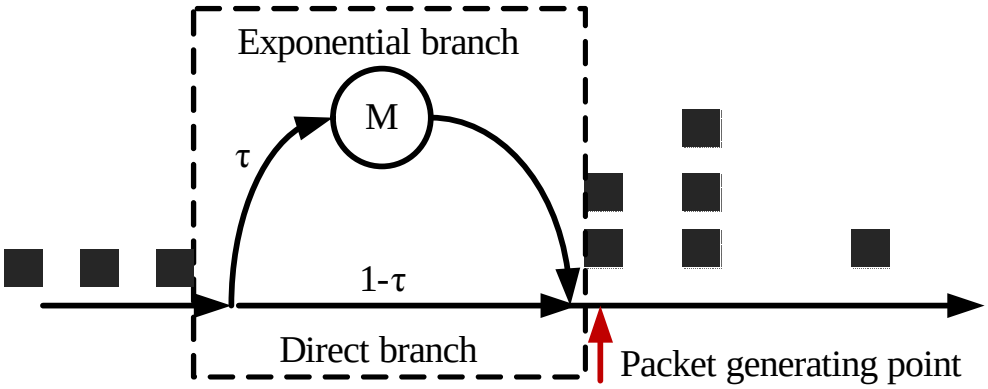

Figure 2.3: Modeling of the GE packet generation process in [7]

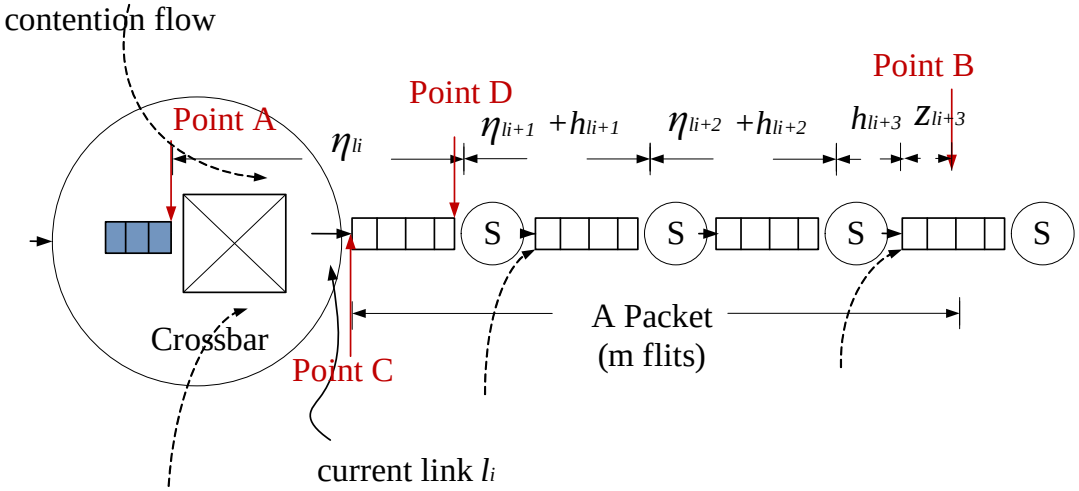

Figure 2.4: An illustration of the channel service time of link l_i under WH routing

with an exponentially distributed service time with mean equal to $1/\lambda\tau$. The burst of packets consist of a packet which comes from the exponential branch (M branch) in addition to a number of continuous packets arriving through the direct branch [7]. The GE distribution is a versatile and simple distribution which helps to make the queuing formulation analytically tractable [7]. Moreover, it has been demonstrated that the GE distribution also provides efficient approximations for short or long-range dependent traffic in supercomputers [7].

2.3.5 WH router modeling

In this section, we present the techniques to estimate the three key components of the WH analytical models, *i.e.,* the flit transfer time η, the path acquisition time h and the source queuing time v_s.

1)*Flit transfer time*: The flit transfer time η of a link l_{ab} is defined as the time

taken for the header flit to leave the buffer head in the upstream node and reach the front of the buffer in the current link l_{ab} after being granted the link access. Fig. 2.4 illustrates this timing concept. It is equivalent to the time taken from Point A to Point D. More specifically, it consists of two parts, the first part is the time to leave the upstream router (*i.e.,* the time from Point A to Point C in Fig. 2.4). For WH routing, this is a constant value depending on the number of pipeline stages (H_s) in the router and the link. The second part accounts for the time the header flit takes to arrive at the front of the buffer of link l_i (q_{l_l}). It is illustrated by the time to travel from Point C to D in Fig. 2.4. This time value can be approximated by the waiting time of a queuing system (such as the M/M/1/K) with capacity equals to the buffer size B. The mean flit arrival rate of this queuing system is given by:

$$\lambda_{flit}^{l_{ab}} = m \times \lambda_{packet}^{l_{ab}} \tag{2.4}$$

$$\lambda_{packet}^{l_{ab}} = \sum_{f \in F_{l_{ab}}} \lambda_f = \sum_{\forall s} \sum_{\forall d} (\lambda_{sd} \times R(s, d, l_{ab})) \tag{2.5}$$

where $R(s, d, k, l_{ab})$ is a binary value indicating whether the channel is an element of the routing path set $P_{s,d}$ [35]:

$$R(s, d, l_{ab}) = \begin{cases} 1 & if\ l_{ab} \in P_{s,d} \\ 0 & otherwise \end{cases} \tag{2.6}$$

The mean time to serve a flit ($s_{flit}^{l_{ab}}$) in this queuing system is calculated by the weighted average of the service time for all flows passing through link l_{ab} and is given by:

$$s_{flit}^{l_{ab}} = \frac{\sum_{\forall f \in F_{l_{ab}}} [\lambda_f \times (\frac{h_{l_{i+1}^f}}{m} + 1)]}{\sum_{\forall f \in F_{l_{ab}}} \lambda_f} \tag{2.7}$$

Eqn. 2.7 shows that for a packet with m flits, it takes $h_{l_{i+1}^f}$ cycles for a header in the buffer to win the next link of flow f. Therefore, $h_{l_{i+1}^f} + 1$ cycles are needed in total for the service of the header flit. After that, the body and tail flits are transferred in 1 cycle without any additional delay. The mean flit service time over the whole packet is thus computed as $h_{l_{i+1}^f}/m + 1$.

After the mean arrival rate ($\lambda_{flit}^{l_{ab}}$) and the mean service time ($s_{flit}^{l_{ab}}$) at this queue are obtained, the M/M/1/K queuing formulation presented in [34] can be applied to approximate q_{l_l}. $\eta_{l_{ab}}$ is then obtained as $\eta_{l_i^f} = q_{l_i} + H_s$.

2)*Path acquisition time*: The path acquisition time h of flow f at its $i-th$ hop l_{ab} (*i.e.,* l_i^f) is defined as the waiting time of the header flit to be granted the access to the buffers in l_{ab} after contention with other flows routing towards the same output direction [36]. It is usually modeled as the waiting time of a queuing system as in [36, 42]. Examples of models used are the G/G/1 [24], M/G/1/K [33, 50], GE/G/1 [7] and MMPP/G/1 [51] queues. For the fair allocation policies such as the round-robin arbitration [52], each flow has the same priority and takes turns to use the output link. Therefore, the system capacity K of the queuing system to derive h is the number of flows that contend for the same link. In Fig 2.4, $k = 3$ for l_i^f. The arrival process of the queuing system is the merging of all flows that route to l_{ab}. More precisely, the mean arrival rate is $\lambda_{l_{ab}} = \sum_{f \in F_{l_{ab}}} \lambda_f$. The higher moments of the arrival process of this queue are calculated according to the specific queuing model used. For example, if the GE/G/1 queuing model is employed to derive h, the squared coefficient of variance (SCV) of the arrival flows are required in order to apply the queuing formula. The SCV of the traffic to l_{ab} can be approximated as:

$$C^2_{a_{l_{ab}}} = \sum_{\forall f \in F_{l_{ab}}} (\lambda_f \times C_f^2) / \sum_{f \in F_{l_{ab}}} \lambda_f \tag{2.8}$$

In [49, 51, 53], detailed mathematical manipulations are shown to obtain a more accurate derivation of $C^2_{a_{l_{ab}}}$. Moreover, the method can be used to generate the estimations of other higher moments that are required by the more generalized queuing models, such as the MMPP/G/1 or MAP/G/1 queues. However, the computation complexity of this method is high as more traffic details are considered. Anyway, it provides a way to model more complex traffic phenomena such as self-similarity and LRD in the queuing models. The service time of the queue that is used to compute h is illustrated in Fig.2.4 (referenced from [42, 49]). Without loss of generality, we assume that a packet has m flits and it spreads over several adjacent links along the path. Here we assume the input buffer depth is B flits. The service time accounts for the time that a packet occupies the link l_{ab} [42]. For example, in Fig. 2.4, assume a header flit in Point A is granted the current link l_i access. Then, the service process of this packet begins when the header flit leaves A and ends when the tail flit departs A so as to release l_i for other flows. If the downstream links are not congested, this service time simply equals to the packet length (*i.e.,* m cycles) because the whole packet can traverse

in a continuous way [42]. However, when there is severe blockage along the path, the worst-case scenario occurs when the packet head reaches Point B (Fig.2.4) and the accumulated buffer spaces from Point C to Point B are just enough to hold the whole packet [42,49]. The time delay for the packet at point A to reach point B, $x_{l_i}^f$ can be given by [42]:

$$x_{l_i}^f = \begin{cases} \eta_{l_i^f} + h_{l_{i+\Lambda_{l_i}^f}} + \sum_{j=1}^{\Lambda_{l_i}^f}(\eta_{l_{i+j}^f} + h_{l_{i+j}^f} + H_s) + z_{l_{i+\Lambda_{l_i^f}}} & if\ \Lambda_{l_i}^f \geq 1 \\ 0 & otherwise \end{cases} \tag{2.9}$$

where $\Lambda_{l_i}^f$ represents the effective number of hops that a packet may span [42, 49]. Specifically, for a flow f with a Manhattan distance from the source to the destination tile equal to d_f, if m/B is smaller than the number of remaining hops $d_f - i$, $\Lambda_{l_i}^f$ equals to $\lfloor m/B \rfloor$; otherwise, it equals to $d_p - i$ [42, 49]. Of note, different from [42], we have included the router pipeline delay H_s in Eqn. 2.9 to reflect a more tight packet transfer time to Point B.

The service time $s_{l_i^f}$ for the flow f is then bounded by m and $x_{l_i}^f$ under different congestion conditions and can be approximated as [42, 49, 50]:

$$s_{l_i^f} = \begin{cases} [m \times (m + x_{l_i}^f) + 2 \times x_{l_i}^f \times m]/(m + 2x_{l_i}^f) & if\ x_{l_i}^f < m \\ [m \times (m + x_{l_i}^f) + 2 \times (x_{l_i}^f)^2]/(m + 2x_{l_i}^f) & otherwise \end{cases} \tag{2.10}$$

The rationale of Eqn. 2.10 is explained as follows [42, 49, 50]: when the downstream channels along the routing path is not congested (*e.g.,* $x_{l_i}^f \to 0$), the link service time $s_{l_i^f}$ is approximated by the packet length m. At the other extreme, when there exists severe blockage at the subsequent hops (*i.e.,* $x_{l_i}^f \to \infty \gg m$), $s_{l_i^f}$ is approximated by $x_{l_i}^f$ as the congestion delay in the subsequent links dominates the current channel service time.

From Eqn. 2.10, once the downstream links transfer time η and contention time h are known, the overall link service time $\bar{s}_{l_{ab}}$ is then the weighted average of the service time of all flows passing through the link l_{ab}, which yields the mean service time $s_{l_{ab}}$ as:

$$\bar{s}_{l_{ab}} = \sum_{\forall f \in F_{l_{ab}}} (\lambda_f \times s_{l_i^f}) / \sum_{\forall f \in F_{l_{ab}}} \lambda_f \tag{2.11}$$

For M/G/1/K, G/G/1 and MMPP/G/1 models, the SCV of the service time for link l_{ab} is also required. It can be approximated in a similar way as [24] and is

given by:

$$C^2_{s_{l_{ab}}} = \frac{\overline{s^2_{l_{ab}}}}{(\bar{s}_{l_{ab}})^2} - 1 = (\frac{\sum_{\forall f \in F_{l_{ab}}} \lambda_f \times s^2_{l^f_i}}{\sum_{\forall f \in F_{l_{ab}}} \lambda_f})/(\bar{s}_{l_{ab}})^2 - 1 \tag{2.12}$$

Based on the above discussions, after obtaining the arrival and the service process characteristics, in this work, we first use the M/G/1/K queuing formula to obtain the waiting time $h^*_{l_{ab}}$ and then extend to $h_{l_{ab}}$ by considering the second moment (SCVs) of the GE traffic input of the queuing system:

$$h_{l_{ab}} = \frac{(C^2_s + C^2_a)}{(1 + C^2_s)} \times h^*_{l_{ab}} \tag{2.13}$$

where $h^*_{l_{ab}}$ is the M/G/1/K queue based waiting time, C^2_s and C^2_a are the SCVs of the service and inter-arrival time, respectively.

3) *Source* queuing time: The source queue at the network interface (NI) is modeled as a queuing system with infinite capacity, and hence an M/M/1/∞ [36] or a GE/G/1/∞ [7] queuing model can be used. For example, if GE/G/1 model is used, the queuing time is given by [7]:

$$v_s = \frac{\bar{s_{l_s}}}{2}(1 + \frac{C^2_a + \lambda_a \times \frac{(\bar{s_{l_s}} - m)^2}{\bar{s_{l_s}}}}{1 - \lambda_a \times \bar{s_{l_s}}}) - \bar{s_{l_s}} \tag{2.14}$$

where the arrival process (λ_a,C^2_a) and the mean service time $\bar{s_{l_s}}$ at the source node s (link l_s) are calculated in similar manners as presented for the router channels.

2.3.6 Discussion on extensions to VC router modeling

When there are multiple virtual channels sharing a single physical port, the inputs received from a physical link will be multiplexed among all the available VCs [36, 55]. In order to model the effect of VC multiplexing, we need to scale the mean packet waiting time by a factor $\bar{V}$, which represents the average degree of VC multiplexing that takes place at a given link channel as in [7, 54–56]. The mean message delay for the flow $f_{s,d}$ is rewritten as:

$$L_{s,d} = (v_s + \eta_{s,d} + h_{s,d}) \times \bar{V} \tag{2.15}$$

To obtain the value of $\bar{V}$, the VC state transition diagram (STD) of a physical channel is used for the analysis [54]. As shown in Fig. 2.5, state V_i $(0 \leq i \leq V)$

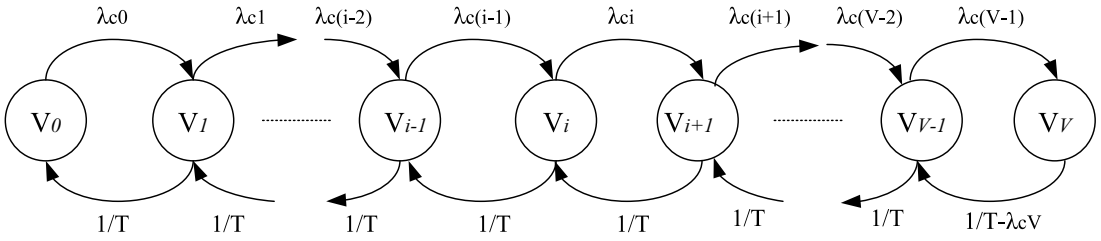

Figure 2.5: The transition diagram to calculate the average degree of virtual channel multiplexing [7]

in the STD represents that there are i VCs currently being occupied by packets in a physical channel. The transition rate from state V_i to V_{i+1} is denoted as λ_{ci}, which reflects the mean packet arrival rate at the given input channel when the STD is at state V_i [54]. It not only depends on the current state (i.e., the number of VCs being used) in STD, but also is related to the traffic arrival model. In [7] and [51], λ_{ci} is derived based on the maximum entropy principles [57] for the GE and MMPP traffic arrival models. Similarly, the rate from state V_i to V_{i-1} is $1/T$, where T is the mean service time at each physical channel. Based on the STD, the state probability, P_v, $(0 \leq v \leq V)$, that represents v VCs are busy at a given channel, can be determined by solving the steady-state equations of the Markov chain in Fig. 2.5 [7,51]. The average degree of VC multiplexing that takes place at a given physical channel is then given by [55]: $\overline{V} = \frac{\sum_{v=0}^{V} v^2 P_v}{\sum_{v=0}^{V} v P_v}$. Consequently, the source queuing time v_s and the path acquisition time $\eta_{s,d}$ in the VC channel model will equal to those calculated in the wormhole model multiplying $\overline{V}$ as shown in Eqn. 2.15.

2.4 SVR-NoC latency model

To further improve the accuracy of the analytical performance model, we introduce a learning based performance model [58]. The main characteristics of this learning model is that in addition to the using of collected traffic data statistics as training features, we also include the prediction results of the queuing model presented in section 2.3 as part of the features. More precisely, the SVR-NoC model learns and refines the proposed queuing model based on the typical training data from the simulation of the target NoC platform. In this section, we first define two regression functions that need to be learned and then elaborate on the support vector regression techniques for obtaining these two models, respectively.

2.4.1 Channel and source queuing regression models

Based on Eqn. 2.1, we define two queuing delays that make up of the traffic flow latency. The first is the source queuing delay SQ_s which is the waiting time of the packets at the source queue s before injected into the network (*i.e.,* $SQ_s = v_s$). The second is the channel waiting time $CQ_{k,dir}$ which is the total packet transfer time at the direction dir of router k and equals to $h_{k,dir} + \eta_{k,dir}$ in Eqn. 2.1. In this work, we model these two components via two regression functions f_{SQ} and f_{CQ}. We denote the feature sets used in learning f_{SQ} and f_{CQ} by two vectors $X_{CQ} = [x_{cq1}, x_{cq2}, \cdots, x_{cqn}]$ and $X_{SQ} = [x_{sq1}, x_{sq2}, \cdots, x_{sqn}]$, respectively. The proposed supervised SVR-NoC learning framework is applied on the training data set to formulate the following models in which the estimated queuing delays are specified as functions of the selected features:

$$\widehat{CQ}_{k,dir} = f_{CQ}(X_{CQ}); \quad \widehat{SQ_s} = f_{SQ}(X_{SQ}) \tag{2.16}$$

Using Eqn. 2.16, given the feature sets, we can estimate SQ_s and $CQ_{k,dir}$ and substitute them in Eqn. 2.1 to obtain the latency of each specific flow. Once the $CQ_{k,dir}$ metric is estimated, the channel buffer utilization and the network throughput can also be obtained as in [35].

2.4.2 Overall SVR-NoC methodology

In SVR-NoC, the step to obtain the estimates of the channel and source queuing delay using the proposed PQ model is summarized in Algorithm 2. In the training stage, we first apply the link dependency analysis presented in section 2.3.3 to determine the correct link ordering for performance analysis. Then, for each link l_{ab} in the ordered list G, we calculate the arrival traffic model $(\lambda_{l_{ab}}, C^2_{al_{ab}})$ according to the routing algorithm and applications. As shown in Eqn. 2.11, the link transfer time $\eta_{l_{ab}}$ only depends on its downstream link contention time $h_{l^f_{i+1}}$, which should have been analyzed in the previous loop. On the other hand, the path acquisition time $h_{l_{ab}}$depends not only on its downstream link contention time $h_{l^f_{i+1}}$ and transfer time $\eta_{l^f_{i+1}}$ but also on the current link's $\eta_{l^f_i}$ (Eqn. 2.9 and Eqn. 2.10). Therefore, $\eta_{l_{ab}}$ is calculated first. After $\eta_{l_{ab}}$ is obtained, we then calculate the mean and SCV of the link service time $(s^-_{l_{ab}}, C^2_{sl_{ab}})$. If this link

connects two routers, we utilize GE/G/1/K queue to obtain h_{lab}. Otherwise, we calculate the source queuing time v_a. Once all the (h, v, η) variables are known, the predicted channel waiting time CQ^* can be obtained as $h+\eta$ while the source waiting time $SQ^* = v$.

Algorithm 2 Analytical working flow to obtain $CQ^*$and SQ^*

Input: F the application flow set; G the ordered list of links for queuing analysis
Output: $CQ^*_{l_{ab}}$ the estimated channel waiting time for link l_{ab}; $SQ^*_a$the estimated source waiting time
for all link $l_{ab} \in G$ **do**
  $(\lambda_{l_{ab}}, C^2_{al_{ab}}) = \text{traffic\_model}(F_{l_{ab}})$
  $\eta_{l_{ab}} = calculate\_transfer\_time(\lambda_{l_{ab}}, s^{l_{ab}}_{flit}, m, B)$;
  for all $f \in F_{l_{ab}}$and $l^f_i = l_{ab}$ **do**
    $s_{l^f_i} = calculate\_link\_service\_time()$;
  end for
  $(s^-_{l_{ab}}, C^2_{sl_{ab}})$=$service\_time()$;
  if $a \neq b$ **then**
    $h_{l_{ab}} = GE\_G\_1\_K\_queue(\lambda_{l_{ab}}, C^2_{al_{ab}}, s^-_{l_{ab}}, C^2_{sl_{ab}}, k)$;
    $CQ^*_{l_{ab}} = h_{l_{ab}} + \eta_{l_{ab}}$
  else
    $v_a = GE\_G\_1\_queue(\lambda_{l_{ab}}, C^2_{al_{ab}}, s^-_{l_{ab}}, C^2_{sl_{ab}},)$;
    $SQ^*_a = v_a$;
  end if
end for

Figure 2.6 shows the overall SVR-NoC framework, which consists of two parts, namely the *training stage* and *prediction engine*. In the training stage, different traffic patterns are fed into the system level NoC simulator to collect the statistics of the training data of the channels in the routers and the source queues. The channel queue feature set X_{CQ} includes the packet arrival rates AR and the forwarding probabilities F in the router as well as the calculated channel waiting time CQ^* from the proposed analytical model. The source queue feature set X_{SQ} includes the injection rates λ at the traffic source, the neighboring channel waiting times from the analytic model CQ^* as well as the calculated source queuing time SQ^*. Both X_{CQ} and X_{SQ} are extracted from the simulation results to form the training data set. Support vector regression is then carried out to obtain the f_{CQ} and f_{SQ} models, respectively.

After the training stage, the obtained SVR models are used in the prediction engine to estimate the latency performance for different traffic patterns. The prediction engine can be embedded in the NoC synthesis framework for the inner-loop design evaluation. Of note, the features in X_{CQ} and X_{SQ} are computed from

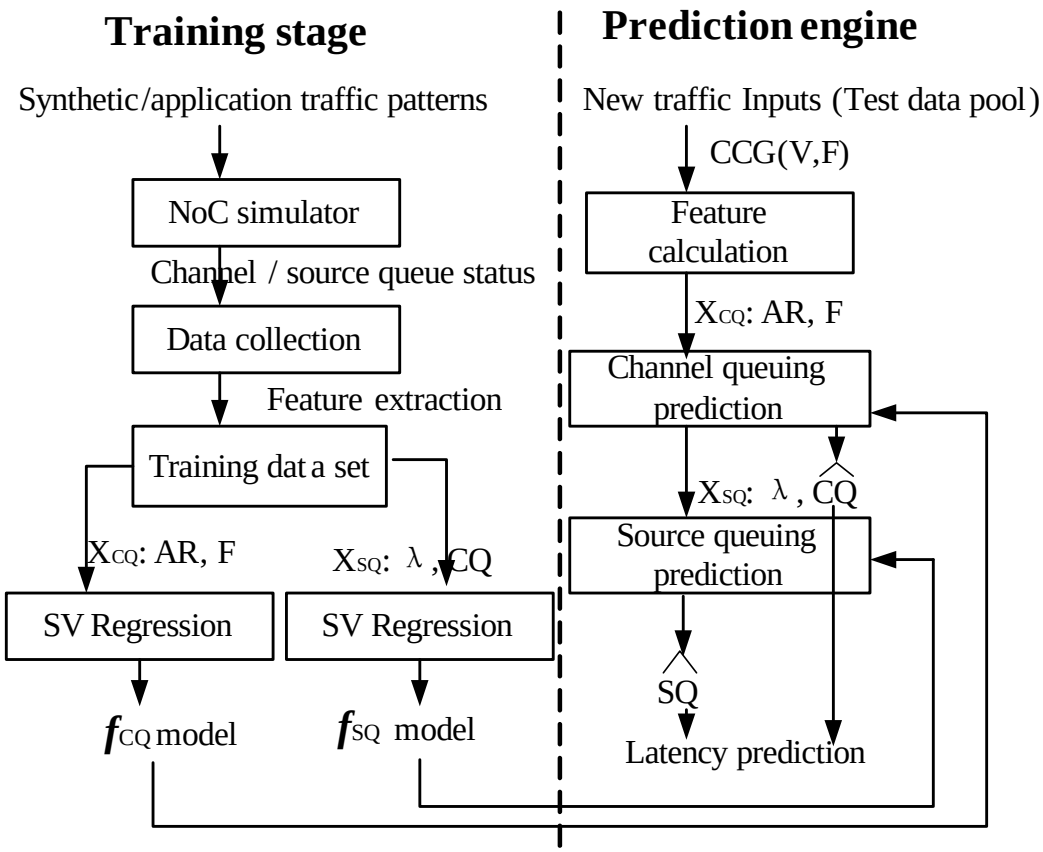

Figure 2.6: SVR-NoC methodology overview

the input core communication graph (CCG) as well as the proposed queuing model. Then, the f_{CQ} and f_{SQ} functions are used to evaluate the channel and source queuing delay, respectively. Finally, the traffic flow latency and the overall latency can be computed according to Eqn. 2.1.

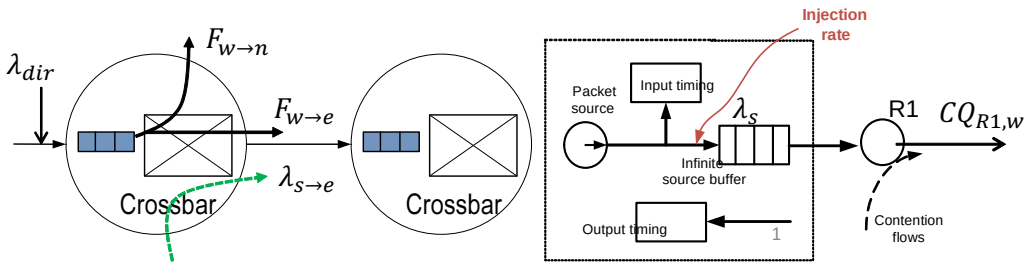

Figure 2.7: Illustration of the channel and source queuing feature vectors

2.4.3 Feature extraction for training data

Channel queuing feature vector

The average waiting time in the channel dir of router k ($CQ_{k,dir}$) depends not only on its packet arrival rate $AR_{k,dir}$ but also on the contention among the channels within the same router, as well as the traffic conditions in the neighboring routers (such as the congestion in the downstream router). In the analytical models, these contentions are usually calculated explicitly by computing the contention matrix using the application traffic arrival rate matrices and the forwarding probability matrices [33, 35]. In contrast, we aim at learning the impact of these contentions implicitly by providing the SVR engine with sufficient samples. Therefore, for the channel queuing feature vector, it is made up of three parts: 1)we first include the arrival rate λ_{dir} of the current channel to indicate the traffic workload injected to the channel. 2)The forwarding probability from current channel to different output directions F_k and the amount of traffic from other input channels to the same output direction λ_k ($k \neq dir$) are included to reflect the contentions situations. 3) Finally, the estimated channel service time S_{dir} and the channel waiting time $\eta_{dir} + h_{dir}$ from queuing models are included to provide an estimation for CQ. Of note, both the arrival rate matrix AR_k and the forwarding probability matrix F_k at node k can be calculated offline based on the application CCG. Specifically, as in [35], $AR_{k,dir}$ is given by

$$AR_{k,dir} = \sum_{\forall s \in V} \sum_{\forall d \in V} (\lambda_{s,d} \times R(s,d,k,dir)) \tag{2.17}$$

where $R(s,d,k,dir)$ indicates whether the channel is part of the routing path set $P_{s,d}$ and is equal to (as in [35]):

$$R(s,d,k,dir) = \begin{cases} 1 & if \quad (k,dir) \in P_{s,d} \\ 0 & otherwise \end{cases} \tag{2.18}$$

The forwarding probability $F_{k,i,j}$ denotes the portion of traffic that arrives at channel i of node k and is forwarded to the output direction j. It is given by (as in [35]):

$$F_{k,i,j} = \frac{\sum_{\forall s \in V} \sum_{\forall d \in V} (\lambda_{s,d} \times g(s,d,k,i,j))}{AR_{k,i}} \tag{2.19}$$

Table 2.3: Training features for regression

Feature vector	Elements notation	Elements	Description
X_{CQ}	AR_{dir}	$X_{CQ}[1]$	The arrival rate in current channel dir
	AR_k	$X_{CQ}[2:5]$	The aggregated traffic from other input channels in router and contends with dir in the output direction k
	$F_{i,j}$	$X_{CQ}[6:9]$	Forwarding probability from current channel i to four output directions
	CQ^*	$X_{CQ}[10]$	The estimated channel queuing time $\eta_{dir} + h_{dir}$
	S_{dir}	$X_{CQ}[11]$	The estimated channel service time
X_{SQ}	Λ_s	$X_{SQ}[1]$	Arrival rate at source s
	F_s	$X_{SQ}[2:5]$	Forwarding probability from s to different output directions
	$lamda_k$	$X_{SQ}[6:9]$	The aggregated traffic that contends with s to the output direction k
	CQ_k^*	$X_{SQ}[10:13]$	The predicted waiting time in the downstream channel of output direction k
	s_k^*	$X_{SQ}[14:17]$	The predicted channel service time in the downstream links of output direction k
	SQ^*	$X_{SQ}[18]$	The estimated source queuing time from the proposed queuing model
	S^*	$X_{SQ}[19]$	The estimated source service time from the proposed queuing model

where $g(s, d, k, i, j)$ is a binary indicator that returns 1 if the routing path $P_{s,d}$ contains the channel (k, i) through the output direction j.

In SVR-NoC, in order to reduce the size of training data that needs to be collected from the simulations, we propose to include an estimation of the channel queuing time CQ^* in the channel queuing feature set. Of note, CQ^* is obtained from the proposed GE/G/1/K queuing model and is independent of the training set simulation results. In Table 2.3, we list all the elements in the channel queuing feature vector X_{CQ}. In Fig. 2.7-a, we show an example of the forwarding probability and aggregated contention traffic. More specifically, as shown in Fig. 2.7, for the current west input channel (dir), when considering the east output direction, the forward probability F_{w-e} represents the portion of traffic that will be forwarded towards east direction. In addition, λ_{s-e} in the figure indicates the traffic that contends with dir to the east output from the south input channel. By including all other input channels, the input feature vector X_{CQ} can be formed as in Table 2.3.

Source queuing feature vector

The time that a packet needs to wait in a particular source queue depends on the traffic generation rate at the processor and the network congestion status, which are reflected by the arrival rate and waiting time in the attached router, respectively. Therefore, the packet generation rate Λ_s is first included in X_{SQ} as shown in Table 2.3. The source generation rate Λ_s can be computed by:

$$\Lambda_s = \sum_{\forall d \in V} \lambda_{sd} \tag{2.20}$$

As shown in Table 2.3, to achieve a better inference on the level of network congestion and a higher accuracy, we also need to include the average waiting time CQ_k of the channels in the downstream links in X_{SQ}, where k is an output direction that the current source queue s has traffic to forward to. Similar to the channel queuing feature set, we include the calculated source queuing time SQ^* in the feature set (shown in Table 2.3) to improve the learning function accuracy. In Fig. 2.7-b, the source traffic rate λ_s as well as the waiting time $CQ_{R1,W}$ in the west down-streaming link are shown, which corresponds to $X_{SQ}[1]$ and $X_{SQ}[10]$ in Table 2.3, respectively.

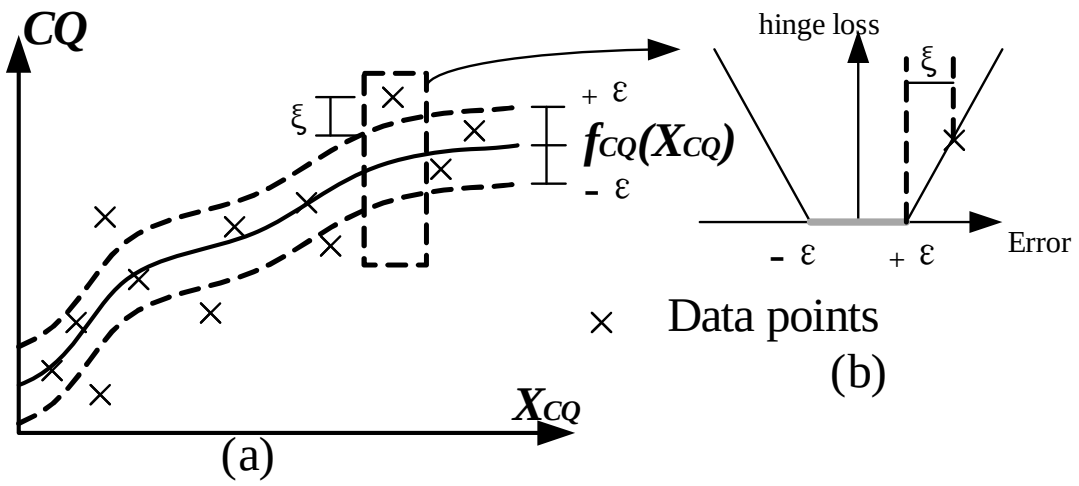

Figure 2.8: The soft margin loss setting in ϵ-SVR [8]

2.4.4 Support vector regression for f_{CQ} and f_{SQ}

After obtaining the feature vectors X_{CQ} and X_{SQ}, we apply ϵ-SVR [38] to learn the two nonlinear models f_{CQ} and f_{SQ}, respectively. The objective of ϵ-SVR is to find a function $f(X)$ that deviates from the actual target values y_i (*i.e.,* $CQ_{k,dir}$ and SQ_s) in the data-set by at most ϵ [38]. There are three steps in the ϵ-SVR [8]: (1) Primal form optimization, (2) Dual problem formulation and (3) Implicit mapping via kernels. The primal formulation is a straightforward representation of the regression problem whereas the latter two steps provide the a practical transformation considering the non-linear extension via kernel tricks. We present the ϵ-SVR formulation of the channel average waiting time function f_{CQ} while similar procedures can be applied for obtaining f_{SQ}.

Primal form formulation

Without loss of generality, we begin by assuming f_{CQ} is a linear function. This assumption will be relaxed later to highly non-linear functions by using the Radial Basis kernels [38] discussed in step 3. Given a set of l training data points, *i.e.,* $\{(X_{CQ1}, CQ_1), \cdots (X_{CQl}, CQ_l)\}$, where $X_{CQi} \in \mathbb{R}^d$ is the i^{th} sample feature vector with a dimensionality of d ($d = 11$ as shown in Table 2.3) . Under the linear model assumption, $\widehat{CQ}$ can be expressed as [8]:

$$\widehat{CQ} = f(X_{CQ}) = \sum_{i=1}^{d} w_i x_{cqi} + b = \mathbf{w^T} X_{\mathrm{CQ}} + b \tag{2.21}$$

where $\mathbf{w} \in \mathbb{R}^d$, $b \in \mathbb{R}$. Let ϵ be the maximum prediction error (*i.e.*, $||CQ - \widehat{CQ}||$ and $|| \bullet ||$ represents the L2 norm) that we can tolerate. We introduce two slack variables, ξ and ξ^*, to represent the error larger and smaller than the target value by more than ϵ, respectively [38]. They are defined as the "soft margin loss". Figure 2.8-a illustrates the soft margin loss for $f_{CQ}(X_{CQ})$. We want to minimize the "soft margin loss" and hence we add penalty if ξ or ξ^* is larger than zero. From Figure 2.8-b, only the points outside the ϵregion will be penalized in a linear fashion [38]. In order to prevent data over-fitting when learning $f(X_{CQ})$ from the training samples, a regularization term proportional to $||\mathbf{w}||^2$ is added into the objective function [8]. To find the optimal f_{CQ}, the problem is formulated as [8]:

$$minimize \quad \frac{1}{2}||\mathbf{w}||^2 + C \times \sum_{i=1}^{l}(\xi_i + \xi_i^*) \tag{2.22}$$

$$subject\ to \begin{cases} (\mathbf{w}^T X_{CQi} + b) - CQ_i \leq \epsilon + \xi_i \\ CQ_i - (\mathbf{w}^T X_{CQi} + b) \leq \epsilon + \xi_i^* \quad (10) \\ \xi_i, \xi_i^* \geq 0, \epsilon \geq 0 \end{cases} \tag{2.23}$$

where $C > 0$ is the regularization parameter that determines the tradeoff between the fitting accuracy over the training samples (the second term in Eqn. 2.22) and the regularization term for preventing data over-fitting (the first term in Eqn. 2.22) [8].

Dual problem expansion

The primal form optimization problem in Eqn. 2.22 and Eqn. 2.23 is difficult to solve, therefore we solve its corresponding Lagrangian formulation $\mathcal{L}$ by introducing a set of Lagrangian multipliers α_i, α_i^*,η_i,η_i^* [8, 38], where:

$$\begin{aligned} \mathcal{L} = \quad & \tfrac{1}{2}||\mathbf{w}||^2 + C \times \textstyle\sum_{i=1}^{l}(\xi_i + \xi_i^*) - \sum_{i=1}^{l}(\eta_i \xi_i + \eta_i^* \xi_i^*) \\ & - \textstyle\sum_{i=1}^{l} \alpha_i \times (\epsilon + \xi_i - CQ_i + \mathbf{w}^T X_{CQi} + b) \\ & - \textstyle\sum_{i=1}^{l} \alpha_i^* \times (\epsilon + \xi_i^* + CQ_i - \mathbf{w}^T X_{CQi} - b) \ (11) \end{aligned} \tag{2.24}$$

It follows that the partial derivatives of $\mathcal{L}$ with respect to the primal variables $(\mathbf{w}, b, \xi_i, \xi_i^*)$ are zero for the optimal point in the primal form [8]. Hence, the primal variables $(\mathbf{w}, b, \xi_i, \xi_i^*)$ can be represented by α_i, α_i^* after setting:

$$\partial_b \mathcal{L} = \partial_{\mathbf{w}} \mathcal{L} = \partial_{\xi_i} \mathcal{L} = \partial_{\xi_i^*} \mathcal{L} = 0 \tag{2.25}$$

By substituting the primal variables in Eqn. 2.22 and Eqn. 2.23 with their dual variable representations, the dual optimization problem can be obtained [38] as follows:

$$max \quad \begin{cases} -\frac{1}{2}\sum_{i,j=1}^{l}(\alpha_i - \alpha_i^*)(\alpha_j - \alpha_j^*) < X_{CQi}, X_{CQj} > \\ -\epsilon\sum_{i=1}^{l}(\alpha_i + \alpha_i^*) + \sum_{i=1}^{l} CQ_i(\alpha_i - \alpha_i^*) \end{cases} \tag{2.26}$$

$$subject\ to\ \sum_{i=1}^{l}(\alpha_i - \alpha_i^*) = 0\ and\ \alpha_i, \alpha_i^* \in [0, C] \tag{2.27}$$

where $< \blacksquare, \blacksquare >$ is the dot product of two vectors. This is a quadratic programming problem and can be solved efficiently in polynomial time. $f_{CQ}(X_{CQ})$ is then obtained as [38]:

$$f_{CQ}(X_{CQ}) = \sum_{i=1}^{m}(\alpha_i - \alpha_i^*) < X_{CQi}, X_{CQ} > +b \tag{2.28}$$

From Eqn. 2.28, it can been seen that the complexity of the regression function f_{CQ} depends only on the training samples X_{CQi} that have nonzero $(\alpha_i - \alpha_i^*)$ terms (which are defined as the support vectors) [8]. Hence the SVR model benefits from keeping only a few training samples for very efficient predictions [8, 38].

Kernel trick for nonlinear extension

In [8, 38], it has been proven that the linear model formulation in Eqn. 2.26 and Eqn. 2.27 is still valid if we substitute the dot product operation $< \blacksquare, \blacksquare >$ with a number of various kernel functions k:

$$k(X_{CQ}, X_{CQ}^{'}) =< \Phi(X_{CQ}), \Phi(X_{CQ}^{'}) > \tag{2.29}$$

From the analytical delay model, it is suggested that the router delay is a non-linear function of the extracted features. Therefore we use the most common Radial Basis Function (RBF) kernel [8] ($k(\mathbf{x}, \mathbf{x}^{'}) = exp(-\gamma \cdot ||\mathbf{x} - \mathbf{x}^{'}||^2)$) in this work, where γ is a tuning parameter to replace the linear dot production operation in Eqn. 2.26. As the RBF kernel is highly nonlinear [8], this kernel trick extends the f_{CQ} formulation from the previous linear assumption into a non-linear function.

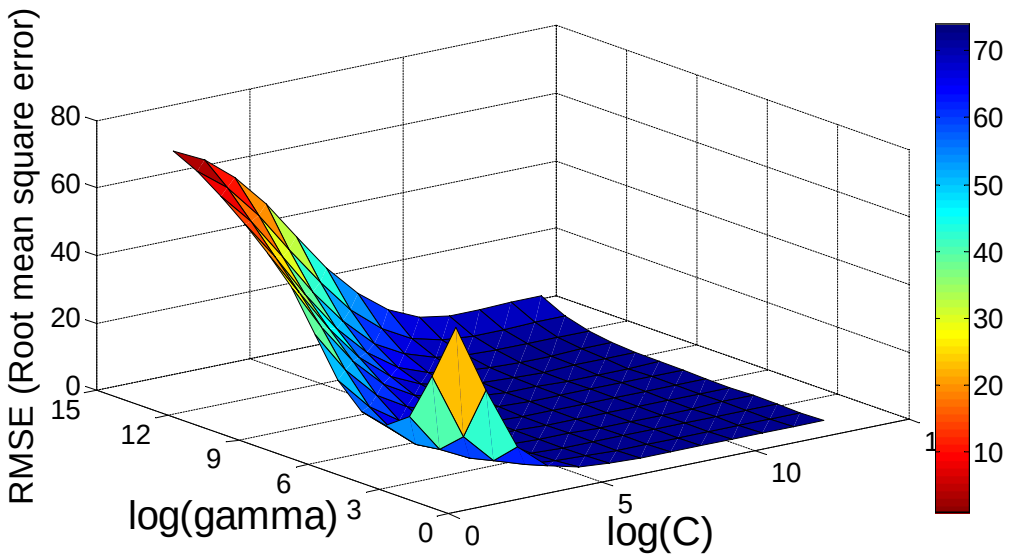

Figure 2.9: An example RMSE surface in the grid search

Cross-validation and grid-search

When solving the optimization problem in Eqn. 2.26-2.27, the tuning parameters ϵ, C, together with the γ in the RBF kernel are assumed to be fixed. They can be set to any arbitrary values. To find out these parameters with highest model accuracy, we need to search over different combinations of (ϵ, C, γ). In this work, we adopt a v-fold cross-validation approach [8, 38]. A v-fold cross-validation separates the training data into v subsets. $v-1$ subsets are used in the training regression stage and the remaining subset is used as the testing set. v runs are carried out with different subsets as the testing set. The cross-validation accuracy is then equal to the root mean square error (RMSE) of all the v runs [8, 38]. In SVR-NoC, we carry out the parameter search on (ϵ, C, γ) using a 5-fold cross-validation as suggested in [59]. Fig. 2.9 shows an example root mean square error surface for ϵ=0.5 with different (γ, C) combinations. In this work, various combinations of (ϵ, C, γ) values are tried and the one with the best cross-validation accuracy (*i.e.,* the lowest RMSE) is selected. The range and resolution of C, γ and ϵ are chosen according to [59, 60].

2.5 Experimental results

2.5.1 Experimental setup and training data preparation

We use Booksim2.0 [61] to simulate the router and NoC performance on mesh NoCs. For each mesh size (ranging from 4×4 to 12×12), various synthetic traffic patterns including *uniform random*, *transpose* and *shuffle* and *tornado* [20] are used as inputs to the target router architecture to obtain the training data-set. To verify the model, after the training stage, the learned model is used to predict the delay with different injection rates for these three patterns. Also, we use this learned model to predict the delay for other traffic patterns that are different from the training sets. Here we use *bitreversal* and *bitcomplement* [20] traffics. The training data set contains runs with different packet injection rates under each traffic pattern, which provides sufficient training data for the SVR learning engine as well as reasonable training time. The SVR-NoC framework is implemented in MATLAB based on Libsvm and LS-SVM library [59,60]. Both synthetic and real application traffic patterns are used in the evaluation. The real application in the experiment includes DVOPD (Dual Object Plane Decoder) [62] which is a video MPSoC benchmark mapped onto a 5×5 mesh NoC.

2.5.2 Proposed queuing model accuracy

We first compare and evaluate the proposed analytical model using two synthetic traffic patterns (*i.e.,*, random and shuffle [20] traffic) with Poisson packet injection rates at each source node. The packet length m is assumed to be fixed of 4 flits. The buffer depth B is adopted to be 9 flits per input port. We compare the proposed model with two state-of-the-art NoC analtyical models. The reference model 1 is adopted from [24], where the priority G/G/1 queue has been modified to a generalized G/G/1 queue formula to model the round-robin arbitration in NoC. The reference model 2 is implemented as proposed in [35] which is based on a generalized M/G/1/ queue. The comparison results of three analytical models against the simulations under 8×8 mesh size are summarized in Fig. 2.10.

As shown in Fig. 2.10, compared to the two reference models, which over- and under-estimate the network saturation injection points by $10\% - 15\%$, the proposed model achieves less than 10% error in predicting the network saturation

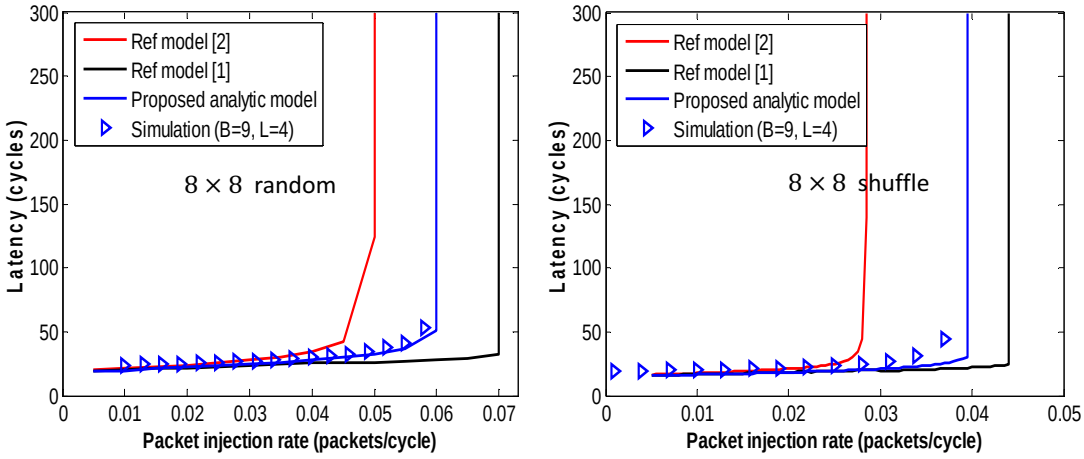

Figure 2.10: Analytical model accuracy comparisons

point (*i.e.,* the injection rate where the overall latency starts increasing dramatically) for both the random and shuffle traffic patterns. Consequently, the proposed analytic model provides a better estimation of the channel and source queuing time and helps to relieve the burden of the learning process.

2.5.3 SVR prediction accuracy

We show the f_{CQ} regression accuracy for the training data-set in Figure 2.11-a. As shown in the figure, our SVR predictor is very accurate for fitting the channel regression curve. The root mean square error (RMSE) between the predicted and the actual values is less than 0.5 and the squared correlation coefficient is higher than 0.99.

Average latency prediction

Although the proposed queueing model provides a comprehensive estimation of different NoC configurations and improves the accuracy under a variety of traffic patterns. From our experimental results, it is still observed the accuracy of the queueing model decreases under some traffic patterns due to the approximations and assumptions made in the derivation. To illustrate this, we first compare the SVR-NoC with the proposed analytical model. In this comparison, 4×4 mesh NoC with random, tornado and shuffle traffic are considered. We consider different buffer depth B and packet length L combinations in the evaluation. The average latencies obtained by different methods are summarized in Fig. 2.11-b,

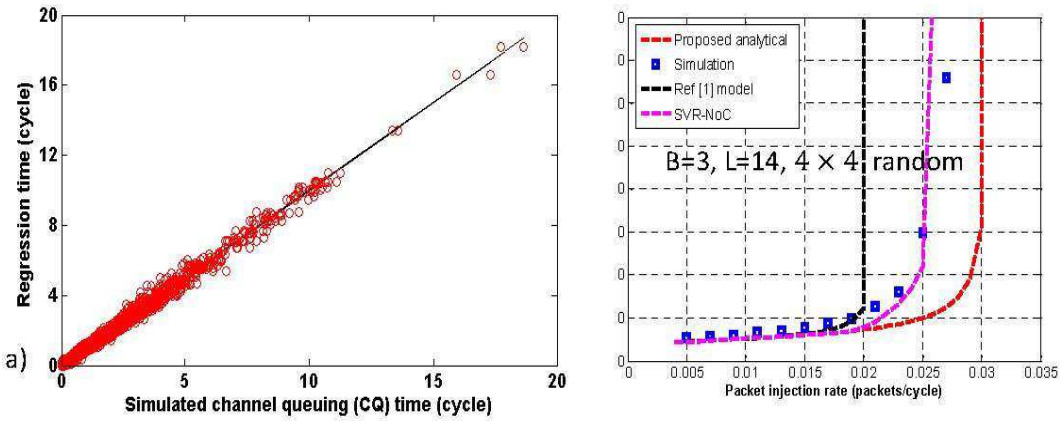

Figure 2.11: a) Regression curve of the channel queuing time (CQ) b)Comparisons of the learning and queuing model

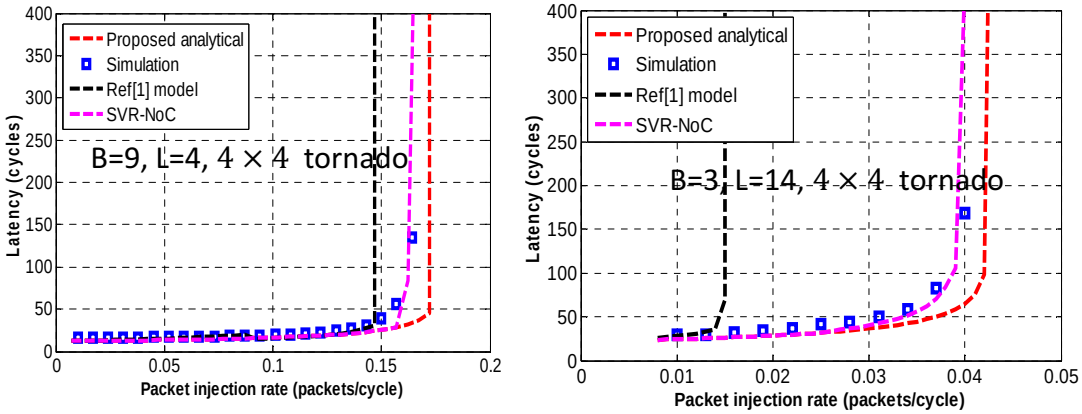

Figure 2.12: Comparisons with the queuing model under different traffic pattern and NoC configurations

Fig. 2.12 and Fig .2.13. As shown in Figure 2.12, the analytical model incurs more than 10% error in predicting the network criticality. Especially for the uneven traffic such as tornado, the error is extremely large for all the packet length and buffer depth combinations. The inaccuracy is due to the assumption and simplification made in the derivation of the PQ model. On the other hand, the SVR-NoC can predict the network saturation point very accurately, with less than 3% error for both the random and transpose traffic patterns.

Figure 2.14-2.16 show the prediction results for various traffic patterns and router architectures. The traffic patterns in these figures are used in the training stage. There are three router architectures considered. They are wormhole routers with: 1) Packet length $L = 4(flits)$ and buffer depth $B = 9$ *flits*; 2) $L = 9$ and $B = 4$ and 3)$L = 14$ and $B = 3$. As can be seen in Figure 2.14-2.16, the errors of SVR-NoC in predicting the network saturation point are within 10% for

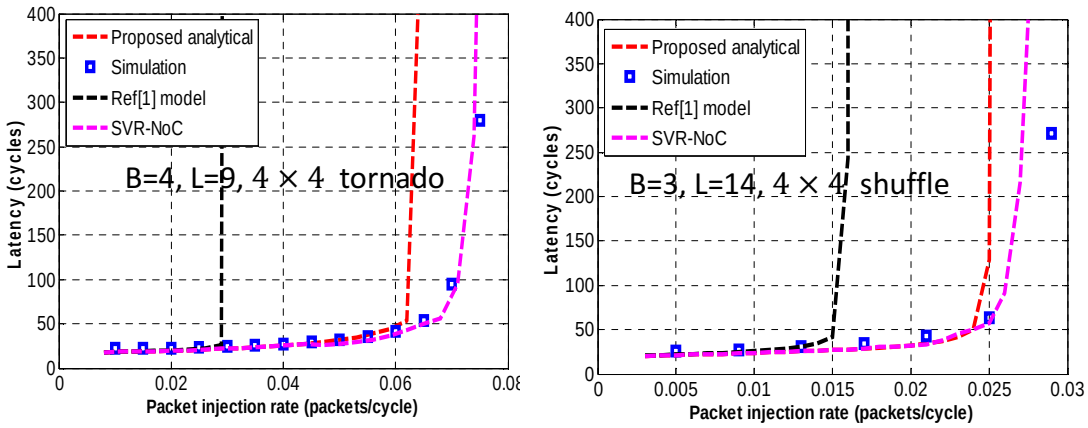

Figure 2.13: Comparisons with the queuing model under different traffic pattern and NoC configurations -2

all the traffic patterns and router architectures.

SVR-NoC for untrained traffic and real applications

Next, we show the latency prediction performance of the SVR-NoC for the other traffic patterns that are not included in the training data set. Figure 2.17-2.18 show the comparison results for 4×4, 8×8 and 12×12 meshes. As shown in the figures, the SVR-NoC still demonstrates high accuracy in the bit-reversal and bit-complement unseen patterns with less than 5% error in predicting the network criticality status. We also use a real benchmark, Dual Object Plane Decoder (DVOPD) [62] for the comparison and demonstrate the accuracy of SVR-NoC. Figure 2.19 shows the comparison of the latency obtained by the simulation and that predicted by SVR-NoC for each traffic flow. As shown in Figure 2.19, the maximum difference between the simulation results and the SVR-NoC regression results is less than 15%. Similar observations can be obtained for several different mappings as well as other application benchmarks.

2.5.4 Runtime comparison

The training time of SVR-NoC for an 8×8 mesh with 5×10^4 samples takes about 6.8 hours. However, the training process is a one-time effort and it does not incur additional overhead during the inner-loop performance prediction. The average simulation time for an 8×8 mesh for 1×10^6 cycles is about fifteen minutes while the SVR-NoC predictive framework takes about eight seconds. The latency estimation method using an analytical model [33] takes about six seconds. We

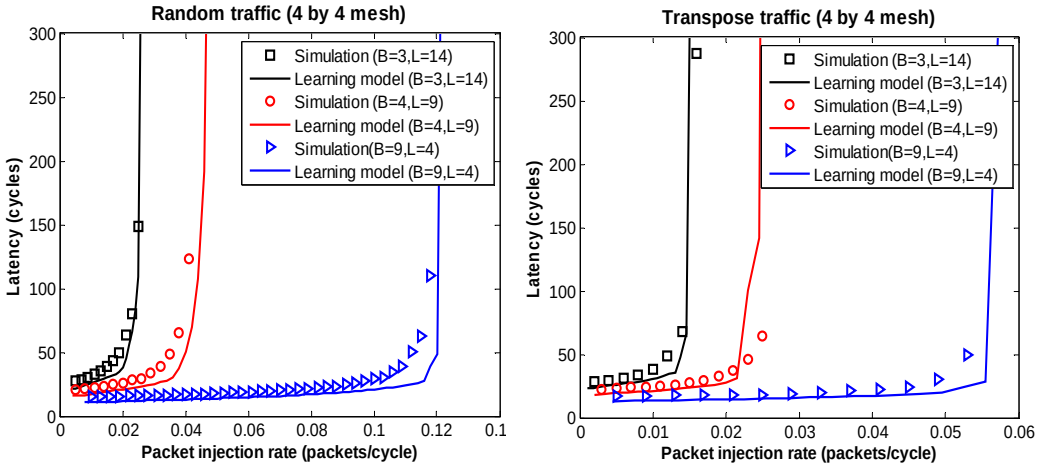

(a) Synthetic traffic used in the training-1

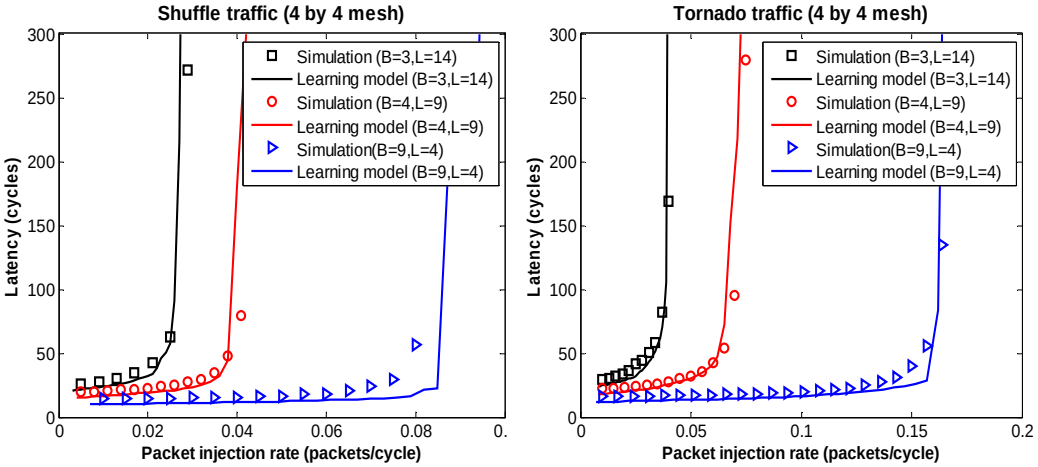

(b) Synthetic traffic used in the training-2

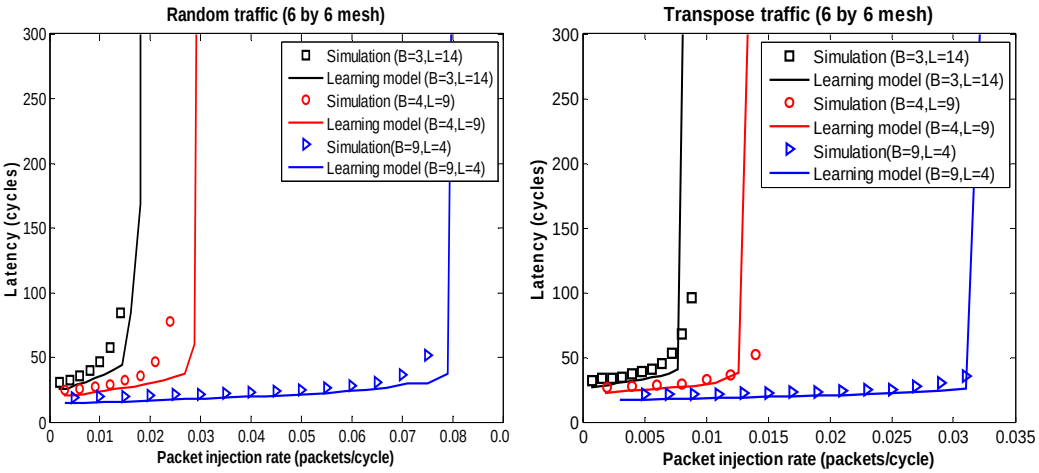

(c) Synthetic traffic used in the training-3

Figure 2.14: Latency prediction performance under trained traffic patterns

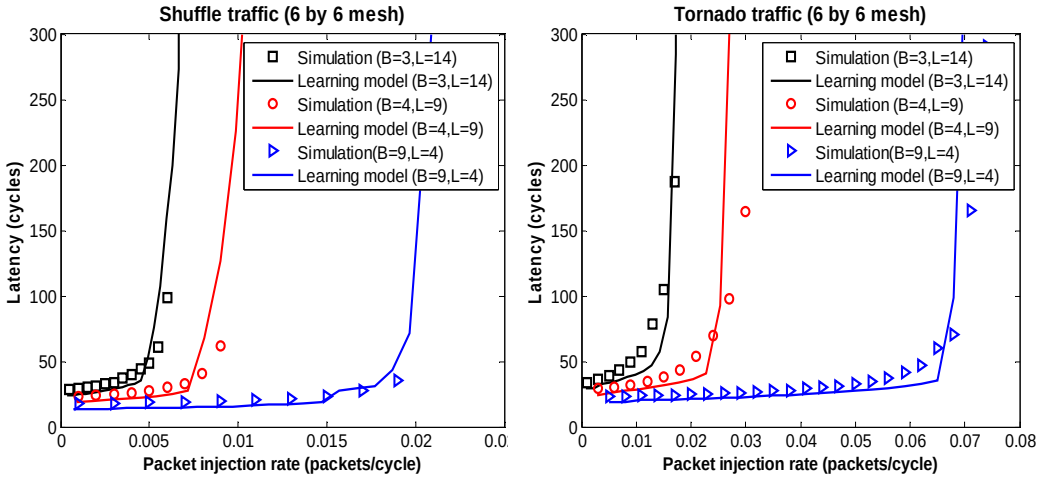

(a) Synthetic traffic used in the training-4

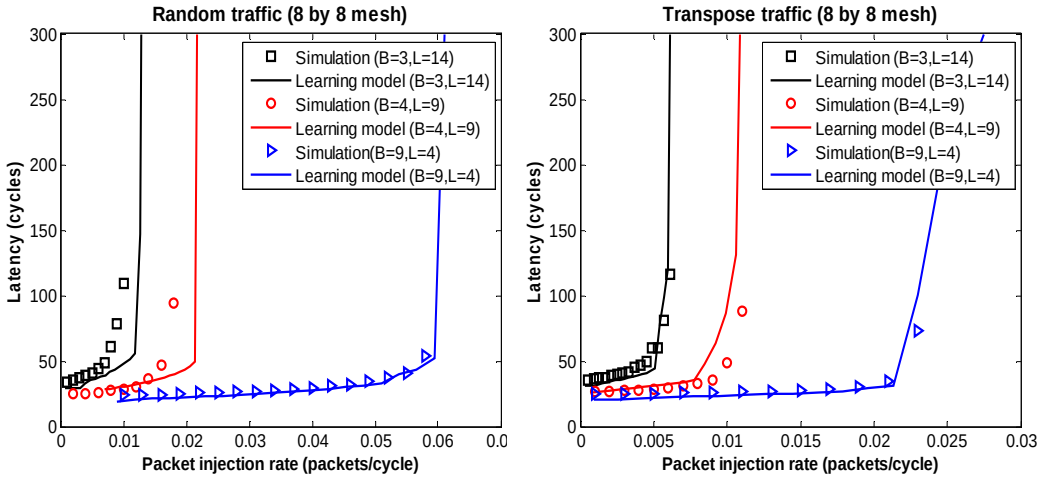

(b) Synthetic traffic used in the training-5

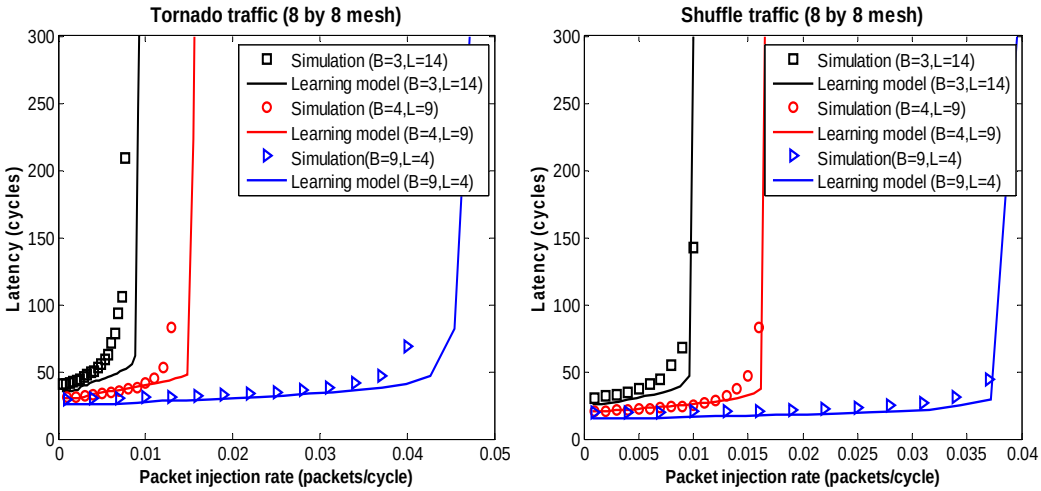

(c) Synthetic traffic used in the training-6

Figure 2.15: Latency prediction for trained synthetic traffics -2

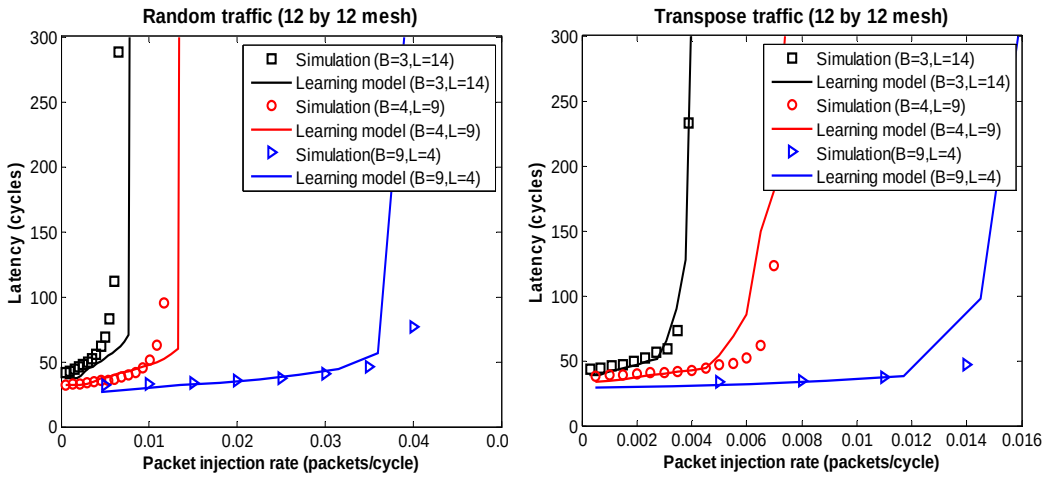

(a) Synthetic traffic used in the training-7

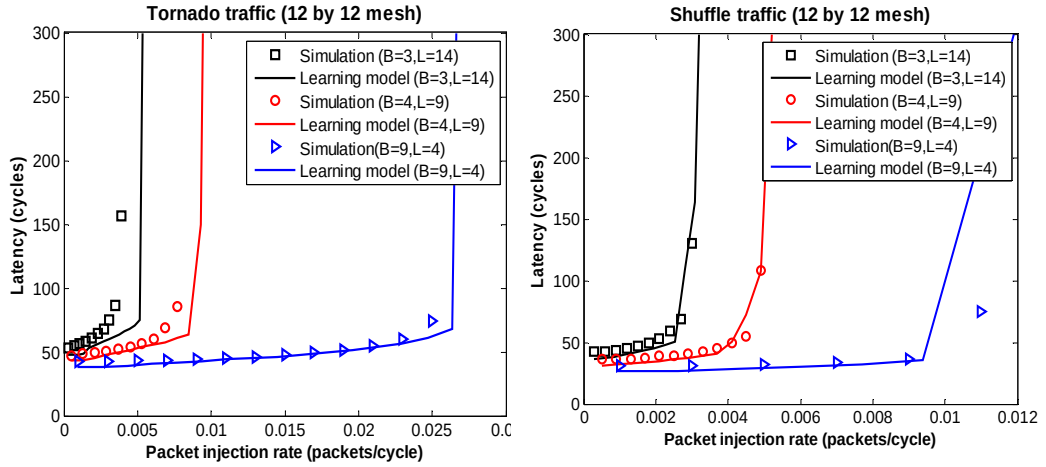

(b) Synthetic traffic used in the training-8

Figure 2.16: Latency prediction for trained synthetic traffics -3

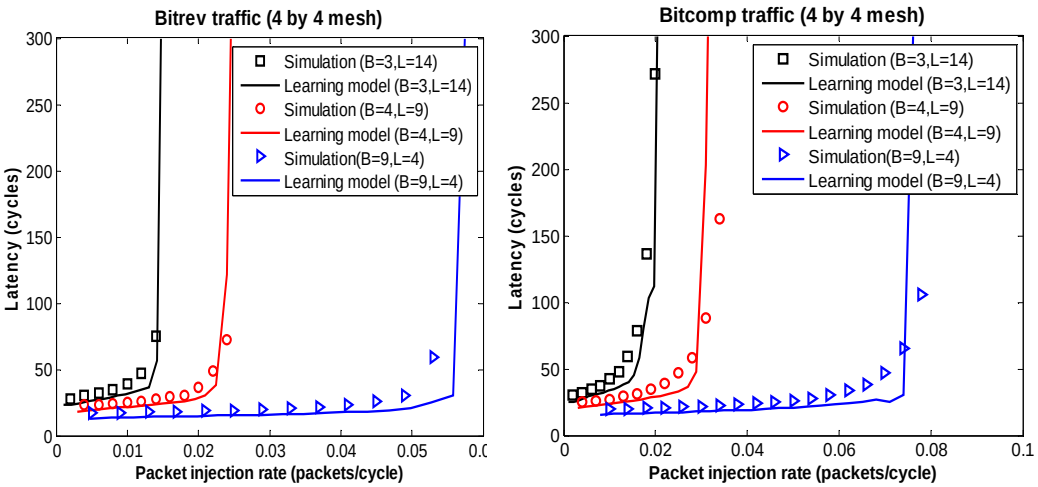

(a) Test the untrained synthetic traffic patterns-1

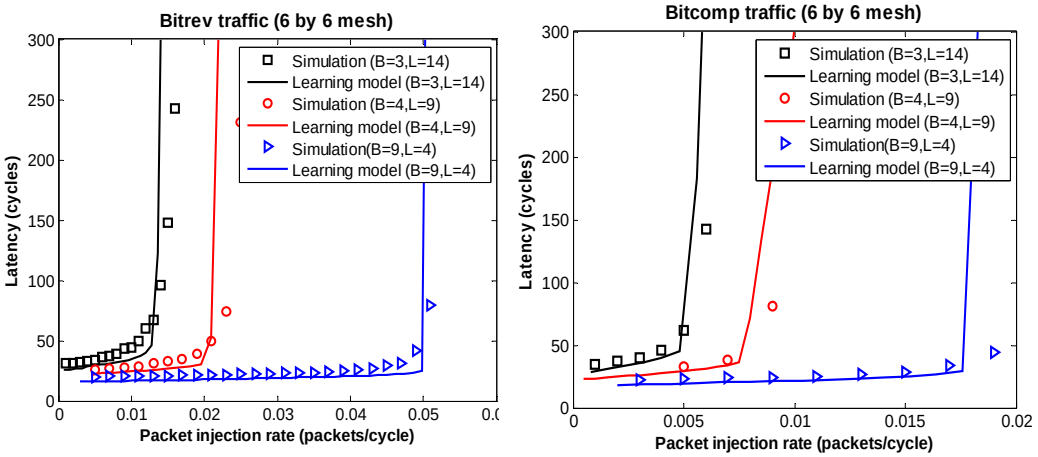

(b) Test the untrained synthetic traffic patterns-2

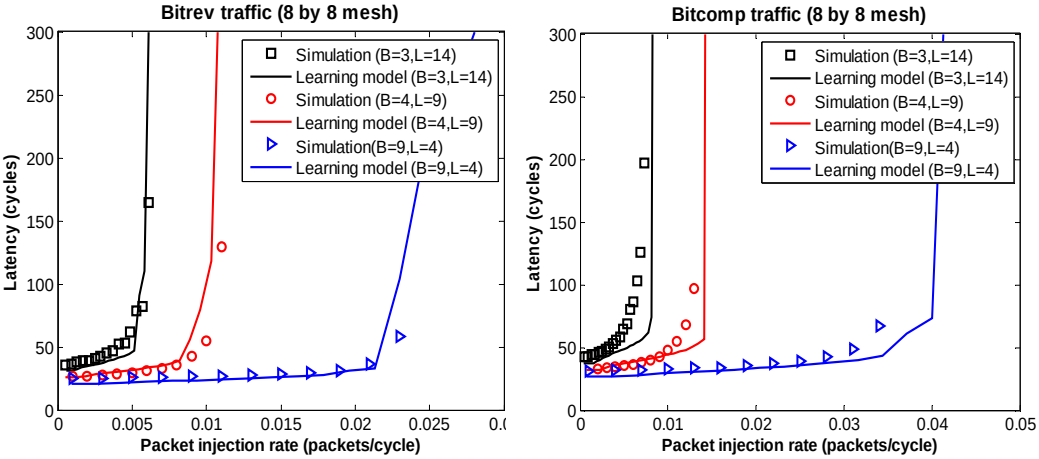

(c) Test the untrained synthetic traffic patterns-3

Figure 2.17: Latency prediction for untrained synthetic traffics

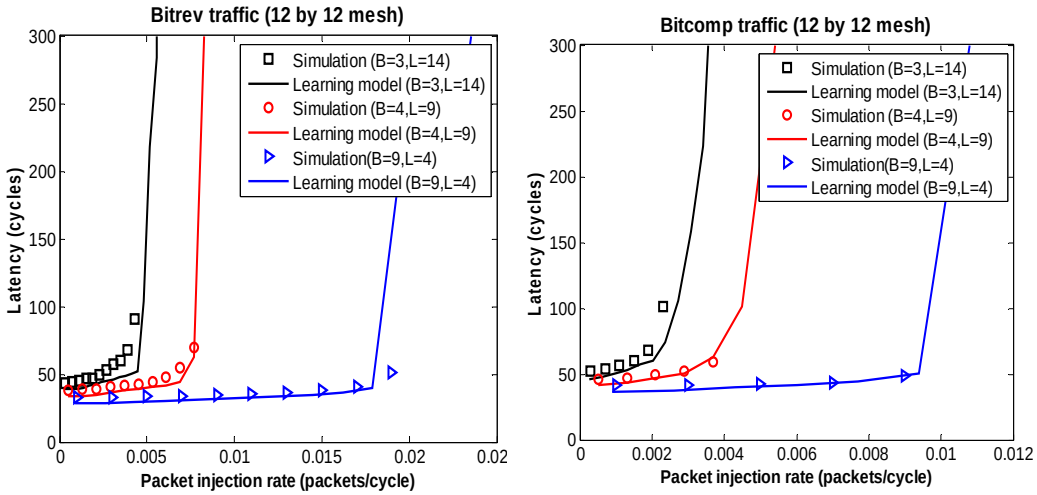

Figure 2.18: Latency prediction for untrained synthetic traffics -2

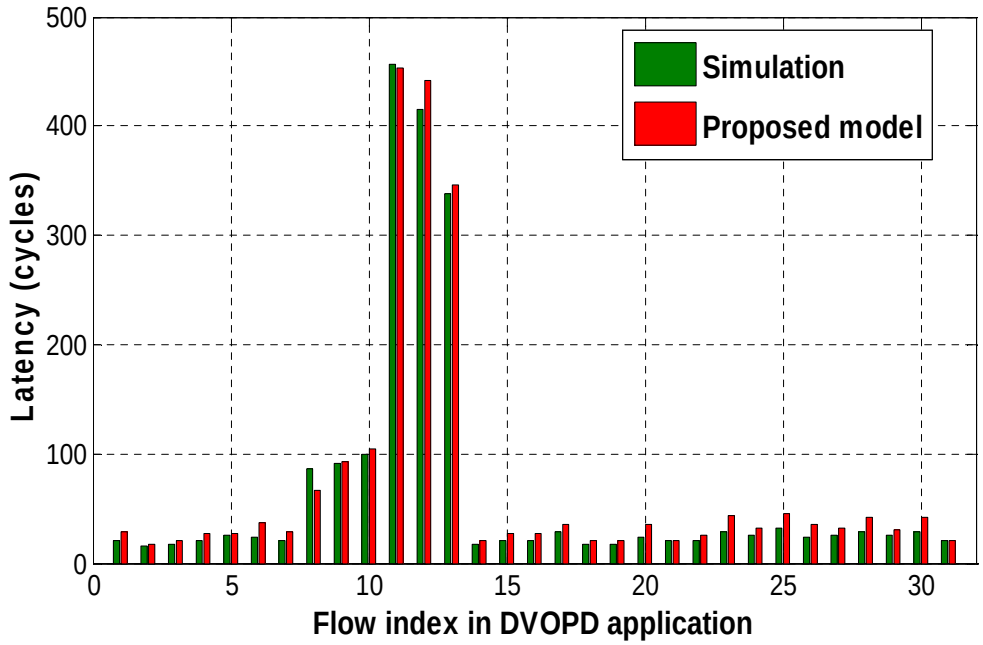

Figure 2.19: Flow latency comparison for DVOPD benchmark

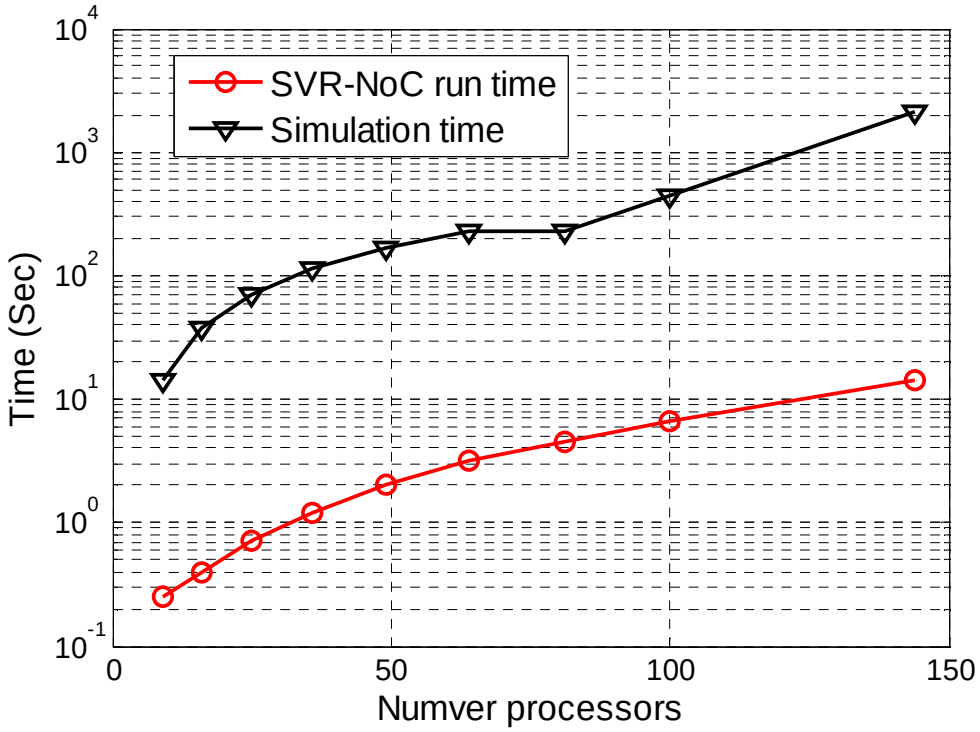

Figure 2.20: Run time comparisons of the learning model and simulation

can see that the SVR-NoC has a similar run-time with the analytical model and has a 120X speedup over the simulation-based prediction. In Fig. 2.20, we show the run time comparisons of the simulations and learning models under different mesh sizes. As can be seen from the figure, the learning model achieves orders of magnitude improvement in the running time, which can benefit the exploration of a large design space in the design phase.

2.6 Conclusion

In this chapter, we propose a machine-learning based approach for NoC performance prediction. We demonstrate the modeling accuracy by using independent test data sets. The SVR-NoC model achieves less than 10% error for various traffic patterns with about 120X speedup compared to the simulation-based estimation.

CHAPTER 3

THERMAL-AWARE AND APPLICATION-SPECIFIC ROUTING ALGORITHM DESIGN

In this chapter, we propose a routing algorithm to reduce the hotspot temperature for application-specific Network-on-chip (NoC). Using the traffic information of the applications as well as the mapping of the tasks onto the NoC, we develop a routing scheme which can achieve a higher adaptivity than the generic ones and at the same time distribute the traffic more uniformly. A set of admissible paths which is deadlock free for all the communications is first obtained. To reduce the hotspot temperature, we find the optimal distribution ratio of the communication traffic among the set of candidate paths. The problem of finding this optimal distribution ratio is formulated as a linear programming (LP) problem and is solved offline. A router microarchitecture which supports our ratio-based selection policy is also proposed. From the simulation results, the peak energy reduction considering the energy consumption of both the processors and routers can be as high as 20% for synthetic traffic and real benchmarks.

3.1 Introduction

For the design of sophisticated MPSoC on high performance NoC, power and temperature have become the dominant constraints [63]. Higher power consumption leads to higher temperature and at the same time, the uneven power consumption distribution across the chip will create thermal hotspots. These thermal hotspots have adverse effect on the carrier mobility, the meantime between failure (MTBF), and also the leakage current of the chip. As a result, they will degrade the performance and reliability dramatically. Consequently, it is highly desirable to have an even power and thermal profile across the chip [64]. This imposes as a design constraint for NoC to avoid uneven power consumption profile so as to reduce the hotspot temperature.

In a typical NoC-based MPSoC design, we need to allocate and schedule the tasks on the available processors and map these processors onto the NoC architecture first. After the task and processor mapping, routing algorithm is developed to decide the physical paths for sending the packets from the sources to the destinations. Each phase of the NoC design affects the total power consumption and also the power profile across the chip. Previously, task mapping and processor core floorplan algorithms [65] have been proposed to achieve a thermal balanced NoC design. However, the routing algorithm is rarely exploited for this purpose. Since communication network (including the routers and the physical links) consumes a significant part of the chip's total power budget (39% of total tile power [66]), the decision on the routing path of the packets will greatly affect the power consumption distribution and hence the overall chip hotspot temperature [63]. Therefore it is important to consider the thermal constraint in the routing phase of the NoC design.

In this chapter, we tackle the routing problem to achieve an even temperature distribution for an application-specific MPSoC. Given an application described by the task flow graph and a target NoC topology, we assume that the tasks are already scheduled and allocated to the processors and the processor mapping is also done. We then utilize the traffic information of the application specified in the task flow graph, which can be obtained through profiling [67, 68], to decide how to split the traffic among different physical routes for an even power profile.

Similar to [69], we use the peak power (or the peak energy under a given time window) metrics to evaluate the effectiveness in reducing the hotspot temperature. Through simulation-based evaluation, we demonstrate that the proposed algorithm can reduce the peak energy of the tiles by $10\% - 20\%$ while improving or maintaining the throughput and latency performance.

3.2 Related Work

In the area of temperature-aware NoC design, many previous works focus on the power consumption distribution of the processor cores. In [70], a dynamic task migration algorithm was proposed to reduce the hotspot temperature due to the processor core (*i.e.,* PE). In [71], a thermal management hardware infrastructure was implemented to adjust the frequency and voltage of the processing elements

according to the temperature requirements at run time.

Since the power consumption of the routers is as significant as the processor core, the thermal constraint should also be addressed in the routing algorithm design. There have been a lot of works on NoC routing algorithms for various purposes including low power routing [71], fault-tolerant routing [19] and congestion avoidance routing to improve latency [72]. However, there are only a few works [63, 69, 73] taking temperature issue into account.

In [69], an ant-colony-based dynamic routing algorithm was proposed to reduce the peak power. Heavy packet traffics are distributed on the chip based on this dynamic routing algorithm to minimize the occurrence of hot spots. However this dynamic routing algorithm is generic in nature and does not take into account the specific traffic information of the applications. Therefore it may not be able to achieve an optimal path distribution. Special control packets are sent among the routers to implement the algorithm which increases the power overhead. Also, two additional forward and backward ant units are needed in the router which results in a large area overhead. Moreover, this work only minimizes the peak power of the routers but does not consider the effect of the processor core power on the temperature. In [63], a run-time thermal-correlation-based routing algorithm is proposed. When the peak temperature of the chip exceeds a threshold, the NoC is under thermal emergency and the dynamic algorithm will throttle the load or re-route the packets using the paths that have the least thermal correlation with the run time hottest regions. However, the algorithm also does not consider the specific traffic information of the applications. It may be inefficient if multiple hotspots occur at the same time. Also it does not clearly describe how to do the re-routing while still guaranteeing the deadlock free property. In [73], a new routing-based traffic migration algorithm VDLAPR and the buffer allocation scheme are proposed to trade-off between the load balanced and the temperature balanced routing for 3D NoCs. In particular, the VDLAPR algorithm is designed for 3D NoCs by distributing the traffic among various layers. For routing within each layer, a thermal-aware routing algorithm such as the one introduced in this chapter is still needed.

In this chapter, we focus on the thermal-aware routing for application specific NoC [74]. To guarantee deadlock free property, generic routing schemes use

algorithms such as X-Y routing, odd-even routing or forbidden turn routing [6]. However, this will limit the flexibility of re-distributing the traffic to achieve an even power consumption profile. Here, we utilize the characterized application traffic information to achieve a larger path set for routing and at the same time provide deadlock avoidance. Higher adaptivity and hence better performance can be achieved since more paths can be used for the re-distribution of the traffic. Given the set of possible routing paths, we formulate the problem of allocating the optimal traffic among all paths as a mathematical programming problem. At run time, the routing decisions will be made distributively according to the calculated traffic splitting ratios. We demonstrate the effectiveness of the proposed routing strategy on peak energy reduction using both synthetic and real application traffics.

3.3 Methodology overview

3.3.1 Assumptions and preliminaries

We aim to achieve an even power consumption distribution and reduce the hotspot temperature for application specific MPSoC using NoC connection. We assume that the given application is specified by a task flow graph which characterizes the communication dependencies and the bandwidth requirements among the tasks of the application. In this chapter, we use a tile-based mesh topology for the NoC due to its popularity and simplicity. However, our methodology can be easily extended to other irregular topologies. Based on the task flow graph and the target topology, task allocation and processor mapping are determined before the routing phase. The energy consumption of each processor core is then estimated according to the tasks mapped to it. Taking all the above information as input, we address the issue of reducing the peak temperature via routing algorithm design. We adopt an adaptive routing strategy instead of static routing scheme because it provides path diversity to distribute traffic among the sources and the destinations [6] and achieves better performance in terms of latency, throughput and congestion avoidance. We also assume minimum path routing is used.

To facilitate the discussion of the proposed methodology and algorithms, we borrow some of the definitions presented and discussed in [23].

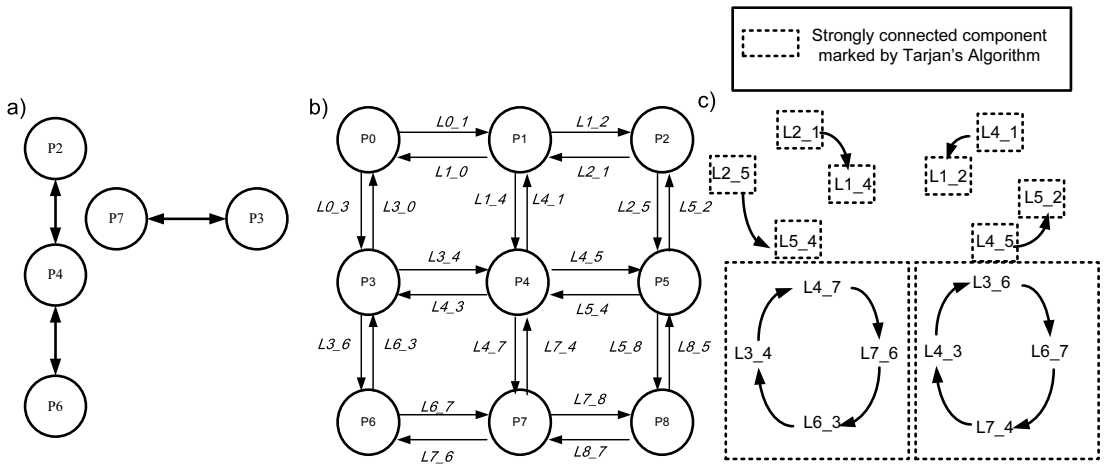

Figure 3.1: An example of: a) Core communication graph (CCG), b) NoC architecture, c) Channel dependency graph (CDG) and its strongly connected components(SCC) after mapping

Definition 3.3.1 *A Task graph* $TG = T(T, C)$ *is a directed graph where the vertex* $t_i \in T$ *represents a task in the application and the edge* $c_{ij} = (t_i, t_j) \in C$ *represents the communication from* t_i *to* t_j*.*

Definition 3.3.2 *A core communication graph* $CCG = (P, E)$ *is a directed graph where the vertex* $p_i \in P$ *represents the* i^{th} *processor in the mesh and the edge* $e_{ij} = (p_i, p_j) \in E$ *represents the communication from* p_i *and* p_j*. Given an application-specific task graph and its corresponding mapping function, CCG can be generated accordingly.*

Definition 3.3.3 *A channel dependency graph* $CDG = (L, R)$ *is a directed graph where the vertex* $L_{mn} \in L$ *represents a physical link* (m, n) *in the mesh topology pointing from tile* m *to tile* n*. An edge* $R_{\mu\nu} = (L_\mu, L_\nu)$ $(\mu = mn, \nu = np) \in R$ *exists if there is a dependency from* L_μ *to* L_ν *which indicates a path passing through the two links* (m, n) *and* (n, p)*consecutively.*

From Duato's theorem [75], a routing algorithm is deadlock free for an application if there are no cycles in its channel dependency graph (CDG).

Definition 3.3.4 *Strongly connected components (SCC): A directed graph is strongly connected if there is a path from each vertex to every other. The strongly connected components of a graph are its maximally connected sub-graphs.*

We illustrate the definitions of CCG, CDG and SCC by an example shown in Figure 3.1. Here we use a 3×3 mesh NoC (Figure 3.1-b as the communication backbone for an application characterized by the core communication

graph shown in Figure 3.1-a. For each communication pair $e_{ij} = (p_i, p_j)$ in the CCG, assume the Manhattan distance (i.e. the minimum number of hops) between the source and destination tile pair (i, j) is $d(i, j) = |i_x - j_x| + |i_y - j_y|$ and the distance between the two tiles in the x-direction is $d_x(i, j) = |i_x - j_x|$, the total number of minimum length paths connecting the two tiles is given by $N_{ij} = C_{d(i,j)}^{d_x(i,j)}$ [1] since we need to traverse along the x direction $d_x(i, j)$ times among total $d(i, j)$ number of hops. Let $l(i, j, k)\,(k \leq N_{ij})$ denote the k^{th} path connecting the two tiles. If $l(i, j, k)$ traverses two network links (m, n) and (n, p) consecutively, then an edge is added to connect the two nodes L_{mn} and L_{np} in the CDG. By inspecting all the feasible paths, we can construct the whole channel dependency graph. In Figure 3.1, taking the communication pair $(2, 4)$ as an example, two minimum length paths exist: $(2 \rightarrow 5 \rightarrow 4)$ and $(2 \rightarrow 1 \rightarrow 4)$.In the CDG, we add two edges connecting (L_{2-5}, L_{5-4}) and (L_{2-1}, L_{1-4}) , respectively. As shown in Figure3.1-c, there are in total 10 strongly connected components in the resulting CDG which can be found by the Tarjan's algorithm [76]. We can find two circles in Figure 3.1-c, *i.e.,* $L_{3-4} \rightarrow L_{4-7} \rightarrow L_{7-6} \rightarrow L_{6-3} \rightarrow L_{3-4}$ and $L_{4-3} \rightarrow L_{3-6} \rightarrow L_{6-7} \rightarrow L_{7-4} \rightarrow L_{4-3}$, where some edges are required to be removed by the path set finding algorithm in Section 4.1 to avoid deadlocks.

Definition 3.3.5 *Routing adaptivity: Routing adaptivity is defined as the ratio of the total number of available paths provided by a routing algorithm to the number of all possible minimum length paths between the source and destination pairs. Higher adaptivity will improve the capability of avoiding congestion and redistributing the traffic to reduce thermal hotspot. However, we need to meet the deadlock-free constraints while improving the routing adaptivity.*

For a given application, we first use the task graph TG (Definition 1) to characterize its traffic patterns as well as the communication volume between the tasks. Then based on the mapping algorithm on the target NoC platform, the task graph TG is transformed to the core communication graph CCG (Definition 2). The channel dependency graph CDG (Definition 3) is built based on the routing paths assigned to each edge in CCG. In this chapter, we propose to find the strongly connected components SCC (Definition 4) of the underlying CDG to remove the circular dependency among the channel resources to avoid potential

[1] C_n^r is the combination notation, where $C_n^r = \frac{n!}{r!(n-r)!}$

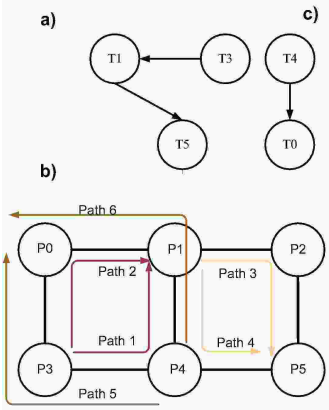

c)

Communication pair	Path used in routing	Number of packets sent over the path	
		Strategy 1 (uniform traffic distribution)	Strategy 2 (allocating traffic non-uniformly)
Tile P3 to P1	Path 1: P3->P4->P1	500	400
	Path 2: P3->P0->P1	500	600
Tile P1 to P5	Path 3: P1->P2->P5	500	850
	Path 4: P1->P4->P5	500	150
Tile P4 to P0	Path 5: P4->P3->P0	500	900
	Path 6: P4->P1->P0	500	100
Router energy distribution		Strategy 1	Strategy 2
Tile 0		1500 E_1	1600 E_1
Tile 1		2500 E_1	2100 E_1
Tile 2		500 E_1	850 E_1
Tile 3		1500 E_1	1900 E_1
Tile 4		2000 E_1	1550 E_1
Tile 5		1000 E_1	1000 E_1
Peak Energy		2500 E_1	2100 E_1
Maximum energy difference		2000 E_1	1250 E_1

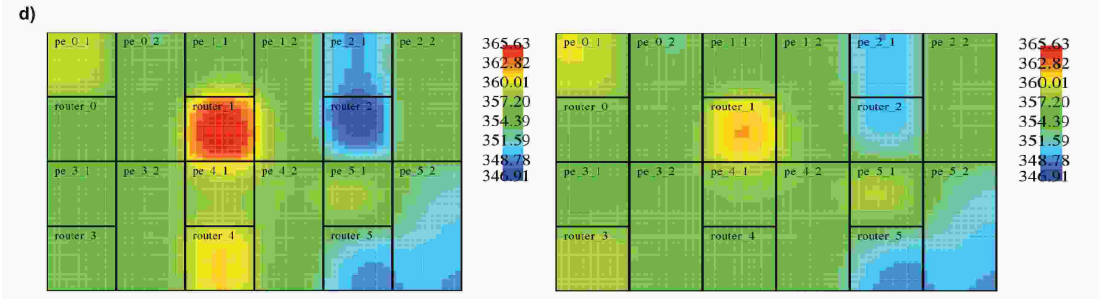

Figure 3.2: A motivation example of allocating traffic among paths: a) the core communication graph(CCG), b) routing paths allocation on 2×3 NoC, c) two strategies of using the routing paths, d) thermal profile comparison (left: strategy one, right: strategy two)

deadlock. Finally, routing adaptivity (Definition 5) is the metric that reflects the capability of redistributing traffic for different cycle breaking algorithms in SCC.

3.3.2 Motivation for thermal-aware routing

For adaptive routing, normally there will be more than one path available for every communication pair. If traffic is distributed equally on all the paths, some of the routers may have more paths passing through them and hence more packets to receive and send. We show the need to allocate traffic properly among the paths by an example illustrated in Figure 3.2. In Figure 3.2, three communication pairs occur concurrently: from tile P3 to P1, tile P1 to P5 and tile P4 to P0. We assume 1000 packets are generated for each pair within a time window and sent over the network in this example. The energy consumption of processing a single packet in a router is denoted as E_1. Two routing strategies are compared. Strategy 1 uses uniform traffic distribution among all paths between the source and the destination nodes for routing. Strategy 2 allocates traffic non-uniformly among the candidate paths. The total number of packets handled by each router is different in these two strategies. The router energy distribution is shown in

Figure 3.2-c. By properly allocating the traffic, we can reduce the peak energy of the tile by 16% and the energy difference among the tiles by 37.5%. In this simple example, we assume the average energy consumption of the processor core and the router are about the same. We use Hotspot [77] to simulate and evaluate the thermal profile. As shown in Figure 3.2-d, strategy 2 indeed makes the thermal profile more uniform and reduces the hotspot temperature.

3.3.3 Thermal-aware routing design flow

From the above example, we can see that we need to find a set of paths for routing the packets for every communication pair and allocate the traffic properly among these paths so as to achieve a uniform power consumption profile. One critical issue of determining the path set is to provide deadlock avoidance. In generic routing algorithms, deadlock is prevented by disallowing various turns [6]. For application specific NoCs, it will be too conservative and unnecessarily prohibit some legitimate paths to be used [23] as some disallowed turns can actually be used because the application does not have traffics interacting with these turns to form circular dependencies in the CDG. By using the information specified by the task flow graph and the derived communication graph, we can find more paths available for routing and increase the flexibility of distributing traffic among the sources and the destinations. Here we use a similar approach as [23] to find the set of admissible paths for each communication pair while still satisfying the deadlock free requirement.

After obtaining the admissible path set for routing, we find the optimal traffic allocation to each path based on the bandwidth requirement to achieve an even power distribution profile. The problem is formulated as a mathematical programming problem and solved by a LP solver. These phases of design are done offline at the design time. After that, the optimal distribution ratio of each path is obtained. For each router and a particular source-destination communication pair, the ratios of the paths passing through it are combined into the probabilities of using each port to route to the destination. At run time, these probability values are stored into the routing table in each router. For each incoming packet, the router will inquiry its routing table and return the candidate output ports for this packet according to the input direction and the destination. The final

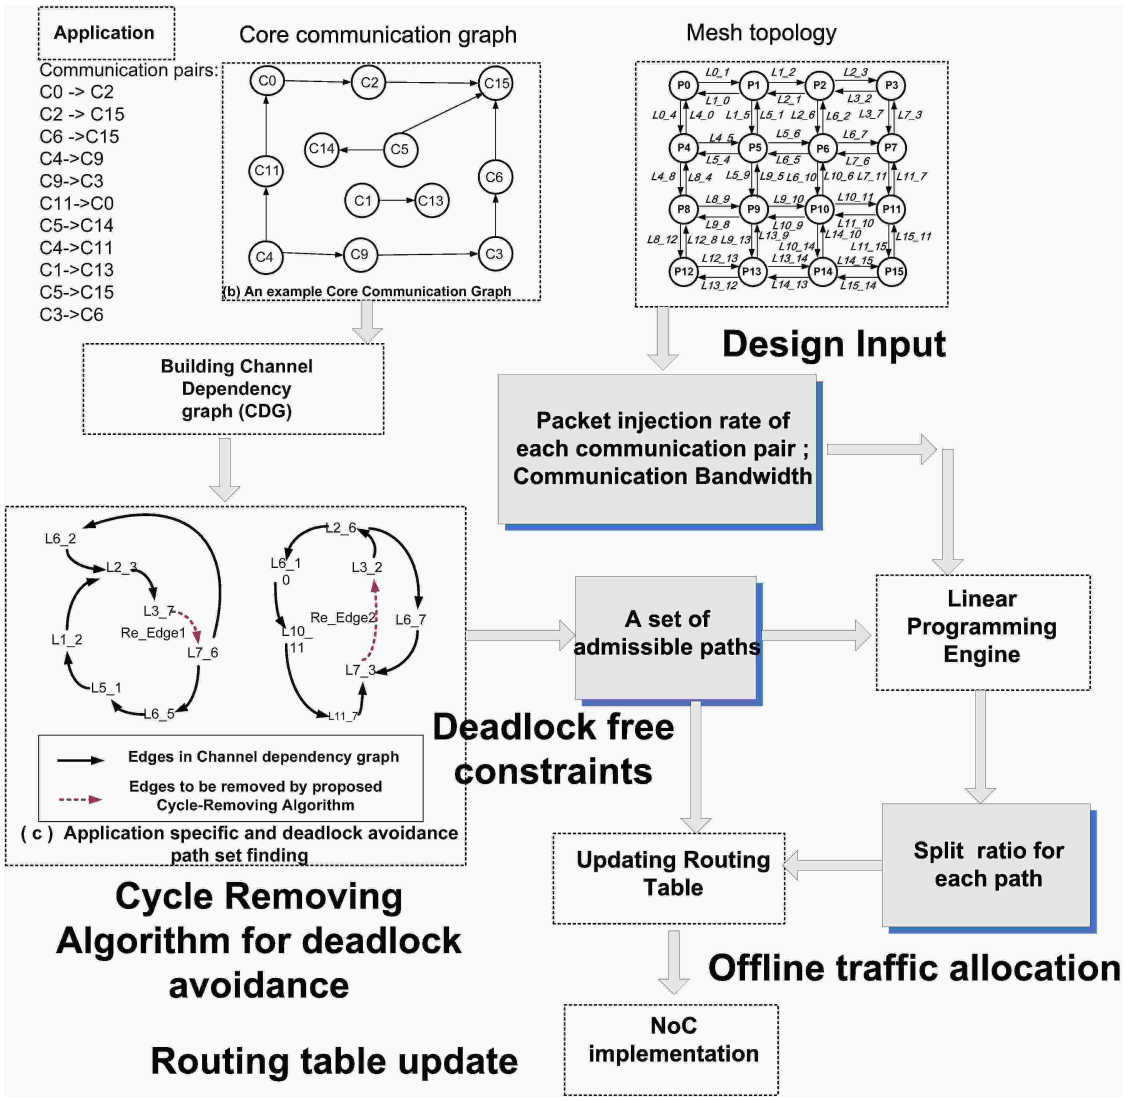

Figure 3.3: Proposed design flow for the thermal-aware routing

output port will be chosen according to the probability value of each candidate.

Figure 3.3 summarizes the whole design flow of our proposed methodology. In the next section, the details of the algorithm will be presented.

3.4 Main algorithm

In this section, we present the details of the offline routing algorithms. We first discuss the algorithm of finding the set of admissible paths for each source-destination communication pair. The admissible paths avoid the circular dependency among any paths and hence provide the deadlock free property. Then we present the optimal traffic allocation problem formulation.

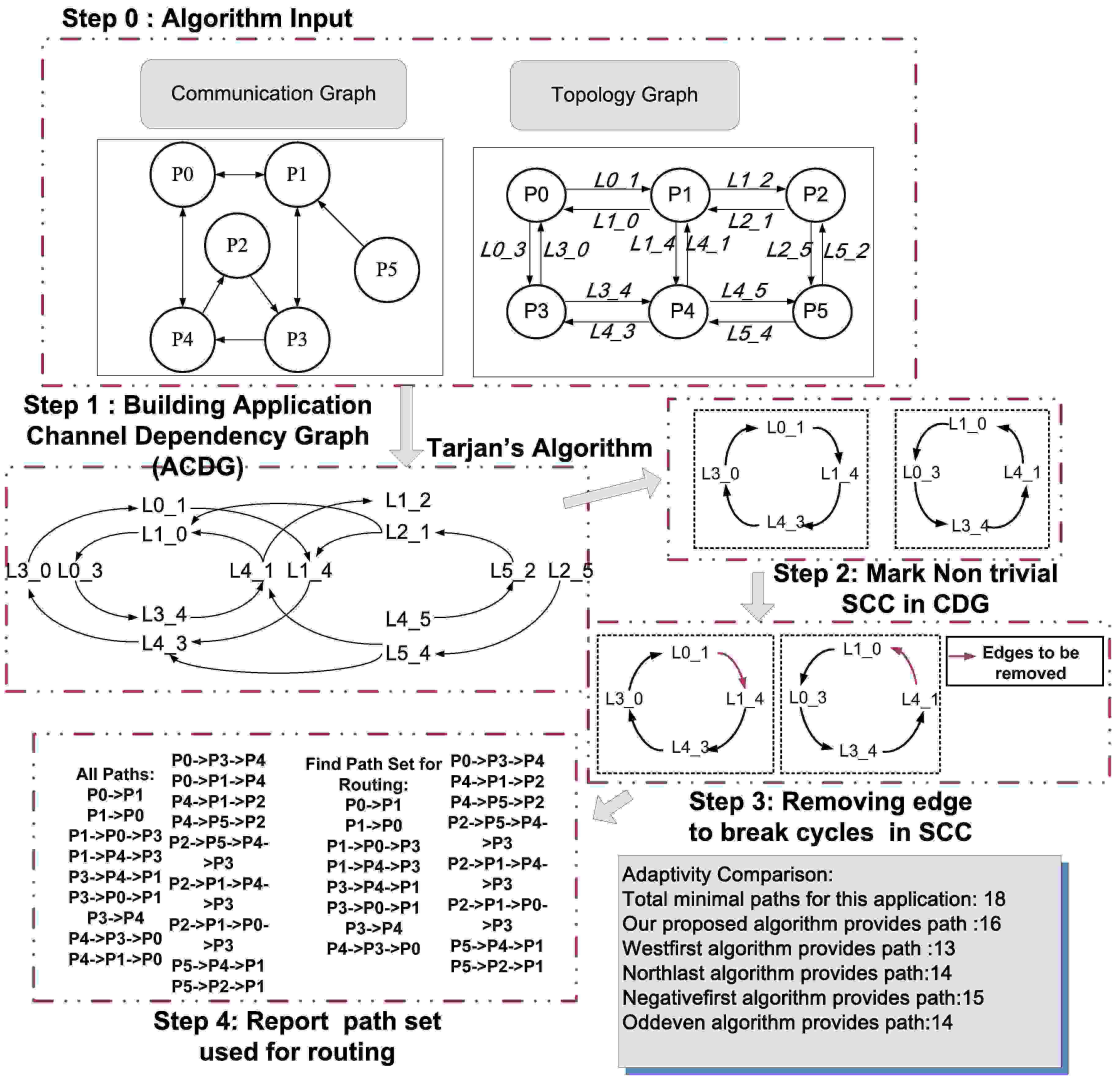

Figure 3.4: Application specific and deadlock free path finding algorithm

3.4.1 Application-specific path set finding algorithm

In [23], a dynamic routing algorithm that increases the average routing adaptivity while maintaining deadlock free is proposed. The average routing adaptivity is often used to represent the degree of adaptiveness and flexibility of a routing algorithm. Here we use a similar approach. Instead of maximizing the average routing adaptivity, we aim to maximize the flexibility to re-divert the traffic to even out the power consumption distribution. Therefore we have to consider the bandwidth requirement of each communication also since the amount of packet processed will directly affect the energy consumption.

Figure 3.4 shows the main flow of our path finding algorithm. Similar to [23], based on the application's task flow graph, we examine all the paths between the

Table 3.1: Notations of application specific path set finding algorithm

Notations	Description
p	a path from tile $src(p)$ to tile $dst(p)$
c	a communication edge in core communication graph
$W(c)$	bandwidth of communication c
C	whole set of communication c in application
$\Phi(c)$	set of minimum paths for communication c
S_{edge}	set of edges to be removed for breaking cycles
$R_{\mu\nu}$	an edge joining two link vertex μ, ν in CDG
$P_{\mu\nu}$	set of paths which introduce $R_{\mu\nu}$in CDG
$\Phi_{S_{edge}}(c)$	set of paths for c after edges in S_{edge} removed

source and destination pairs to build the application channel dependency graph (ACDG). Most likely, there will be cycles in the ACDG so that some edges have to be removed to break these cycles to guarantee deadlock free. In [23], a branch and bound algorithm is used to select the set of edges to remove all the cycles while maintaining all the connectivity and maximizing the average adaptivity. The average adaptivity α is defined as:

$$max\,\alpha = max\,\frac{1}{|C|}\sum_{c\in C}\alpha_c\,;\ \alpha_c = \frac{|\Phi_{S_{edge}}(c)|}{|\Phi(c)|} \tag{3.1}$$

where

$$\Phi_{S_{edge}}(c) = \Phi(c) \setminus \{p|p \in \Phi(c) \bigwedge p \in P_{\mu\nu}\ \forall R_{\mu\nu} \in S_{edge}\} \tag{3.2}$$

The notations used in the above equations are summarized in Table 3.1. Connectivity is guaranteed for every communication c by making sure that at least one path exists between the source and destination nodes. So we have $|\Phi_{S_{edge}}(c)| \geqslant 1;\ \forall c \in C$. In this chapter, instead of finding all cycles in the ACDG and breaking each cycle respectively, as done in [23], we apply Tarjan's algorithm [76] to find all the strongly connected components (SCC) and try to eliminate cycles within each nontrivial components (components containing more than one vertex). One important feature of SCC is that cycles of a directed graph are contained in the same components. We then eliminate cycles within each non-trivial components to achieve deadlock free. Tarjan's algorithm is used because

the complexity is lower ($O(|V| + |E|)$). In many cases, several edges are shared among different cycles (as illustrated in Figure 3.3,the two edges (L_{3-7}, L_{7-6}) and (L_{7-3}, L_{3-2}) are shared among several cycles). If we inspect each cycle separately, we may consider these edges more than once. On the other hand, when we use Tarjan's algorithm, cycles with common edges are in the same component and hence decision can be made more efficiently if we remove some shared edges to break these cycles simultaneously. When we select edges to break the cycle, instead of optimizing the average routing adaptivity, we maximize the following objective function:

$$max \frac{1}{|C|} \sum_{c \in C} \alpha_c W_c = max \frac{1}{|C|} \sum_{c \in C} W(c) \times \frac{|\Phi_{S_{edge}}(c)|}{|\Phi(c)|} \tag{3.3}$$

Here we weight the routing adaptivity of each communication with its corresponding bandwidth requirement ($W(c)$ is the bandwidth of communication c). The rationale is that for the communications that have large bandwidth requirement, more packets need to be processed and routed, the impact on the power consumption distribution is higher. Therefore we should have higher flexibility for these communications to re-divert the traffic and hence higher adaptivity.

3.4.2 Optimal traffic allocation

In the following, we use the notations summarized in Table 3.2 and the energy model described in section 3.4.2 to obtain the linear programming (LP) formulation of the optimal traffic allocation problem.

Energy consumption model

We assume the energy consumption of each processor i (E_{p-i}) is available after task mapping. Wormhole routing is used in our routing scheme. In wormhole routing, each packet is divided into several flits which are the minimum units for data transmission and flow control. For every data packet, the head flit sets up the path directions for the body and the tail flits. Thus, E_{rc}, E_{sel} and E_{vc} only incur when the head flit is processed by the router. Total energy consumption for processing a single packet in router i is given by:

$$\triangle E_{r-i} = (E_{buffer-rw} + E_{forward} + E_{sw}) \times S_{packet} + E_{rc} + E_{sel} + E_{vc} \tag{3.4}$$

Table 3.2: Parameters and notations in LP formulation

Notations	**Network on chip parameters**
N	total number of tiles in NoC
C_{ij}	link capacity between tiles i and j
L_{ij}	number of minimum paths between tiles i and j
$l(i,j,k)$	the k^{th} path joining tiles i and j ; $1 \leqslant k \leqslant L_{ij}$
T_i	set of paths pass tiles i $T_i = \{l(m,n,t) \vert i \in l(m,n,t)\}$
$p(i,j)$	packet injection rate between tiles i and j
S_{packet}	number of flits in a packet
S_{bit}	number of bits in a packet
Notations	**Tile energy parameters**
E_{p-i}	average processor energy of the i^{th} tile
$E_{forward}$	energy of forwarding a flit in the router
$E_{buffer-rw}$	buffer read and write energy for a flit in the router
E_{sw}	switch allocation energy for a flit in the router
E_{rc}/E_{sel}	routing computation/output port selection energy
E_{vc}	virtual channel allocation energy

Let n_i denote the number of packets that received by router i, then the total router energy consumption is equal to $E_{r-i} = \triangle E_{r-i} \times n_i$. The total energy of each tile i (E_i) is equal to $E_{p-i} + E_{r-i}$.

LP problem formulation

Given the set of admission paths for every communication pair which is deadlock free, we want to obtain the ratio of traffic allocated to each path so as to minimize the maximum energy consumption among all the tiles. We formulate the following linear programming problem to solve the optimal path ratios:

1) Variables $r(i,j,k)$**:** ratio of traffic allocated to the k^{th} path $l(i,j,k)$ between tiles i and j among all the L_{ij} paths, where $1 \leq i \leq N_{tile}, 1 \leq j \leq N_{tile}$, $1 \leq k \leq L_{ij}$.

2) Objective functions : The energy consumption of the i^{th} tile, E_i, is

given by

$$E_i = E_{p-i} + \triangle E_{r-i} \times n_i \tag{3.5}$$

where n_i is the total number of packets received by router i per unit time and is equal to the summation of the number of the packets from all paths that pass through tile i , *i.e.,* T_i . n_i is given by:

$$n_i = \sum_{\forall l(a,b,k) \in T_i} r(a,b,k) \times p(a,b) \tag{3.6}$$

In order to balance the tile energy E_i, our objective function is written in a min-max form as follows:

$$obj: \ min(max(E_i)) \ for \ 1 \leq i \leq N_{tile} \tag{3.7}$$

It is equivalent to:

$$min(E) \ s.t. \ E \geq E_i \ for \ 1 \leq i \leq N_{tile} \tag{3.8}$$

3) Problem constraints : The objective function is optimized subject to the following constraints:

3-1) Traffic split constraints : summation of all the traffic allocation ratios between a given pair (i, j) should equal to 1.

$$\sum_{k=1}^{L_{i,j}} r(i,j,k) = 1 \ \forall (i,j) \in C \tag{3.9}$$

$$r(i,j,k) \geq 0 \ \forall i,j \in [1, N_{tile}], \ k \leq L_{ij} \tag{3.10}$$

3-2) Bandwidth constraints : the aggregate bandwidth used for a specific link should not exceed the link capacity.

The communication bandwidth = packet injection rate ×packet size ×clock frequency. Assume T is the cycle time and (i, j) is a physical link in the mesh NoC, a path $l(a, b, k)$ will traverse this link if $l(a, b, k) \in T_i \cap T_j$. So we have

$$\sum_{\forall l(a,b,k) \in T_i \cap T_j} \frac{r(a,b,k) \times p(a,b) \times S_{bit}}{T} \leq C_{ij} \tag{3.11}$$

Given the application's task flow graph and the task mapping, the packet injection rate $p(a, b)$ can be calculated by summing the bandwidth requirement

from all the communication pairs where the source tasks are mapped onto tile a and destination tasks are mapped onto tile b.

The above formulation is a typical linear programming problem and can be solved efficiently using MATLAB CVX optimization toolbox [78].

Using the path ratios in the routers

After solving the LP problem, we obtain a set of admissible paths and their corresponding traffic allocation ratios. To use these information in the implementation of the NoC routing, we can use two schemes. The simplest way is source routing. In the source processor, the header flit of each packet contains the entire path information. The source processor i decides which path to use to send the packet to destination j according to the traffic allocation ratios $r(i, j, k)$ of the set of admissible paths. Intel Teraflops chip [6] uses this source-based routing scheme as the router will be simpler. However, this will create a large overhead on the effective bandwidth as the packet needs to contain additional payload to record the entire routing path.

A more efficient way of implementing the thermal aware routing is to use routing tables stored in each router. One of the major advantages of table-based routing is that it can be dynamically reconfigured or reloaded [23] to allow modifications in the communication requirements. However since the routing decision is made within each router without knowing the entire path information, the path traffic-allocation ratios can not be directly used. In the following, we will show how to convert these ratios into local probability values for the router to select which output port to send out the packet.

To support the thermal-aware routing, the routing table is organized as follows: for each router t in the mesh topology and each of its input direction $d \in D_{in}(t)$, there is a routing table $RT(t, d)$ [23]. For each output port o, there is a set of corresponding entries in the routing table. Each entry consist of the tile id numbers of the source (s), and destination (b) pair, and the probability values $p(o)$ of using o to route to b. Formally, $RT(t, d) = \{(s, b, o, p(o)|o \in D_{out}(t), 0 \leq p(o) \leq 1\}$.

When using routing tables within the routers, the final routing path for a packet of a specific traffic flow is composed by the output port selected in each

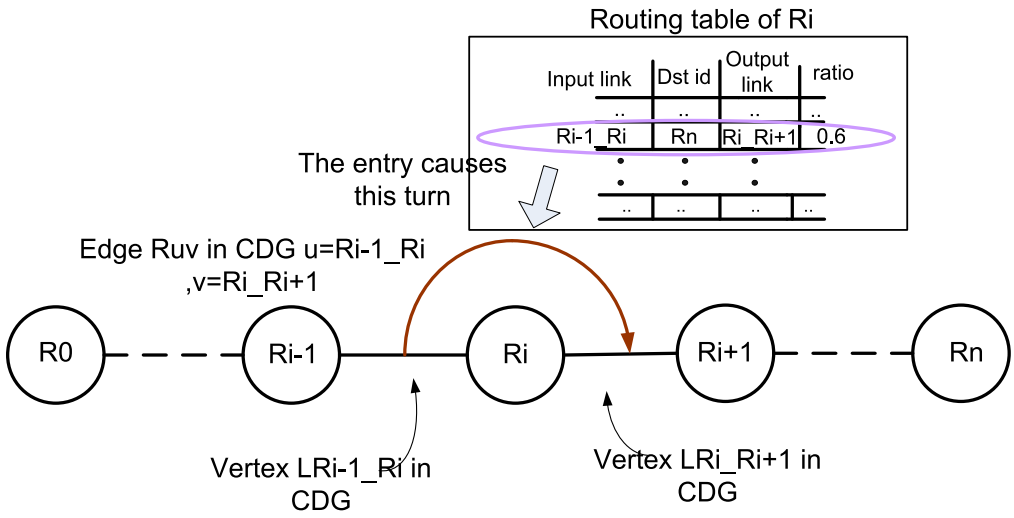

Figure 3.5: Proof of deadlock free inheritance

router along the path distributively. One problem may hereby arise: will the selected output ports in the routers form a new path that is not included in the admissible path set and hence introduce the possibility of deadlock? In the following, we prove that path generated at run-time using the table-based routing is indeed deadlock free provided the path set used for generating the routing table $RT(t, d)$ contains no circles.

Lemma 3.4.1 ***(deadlock-free property in distributive routing):*** *For the table-based routing, after generating the distributed routing table for each router according to the deadlock free path set P, the actual route generated at run time will not contain a disallowed path not included in P. Deadlock-free property is hence inherited from the path set P.*

(by contradiction): We prove that a disallowed path R can not exist if all the entries of the routing table are created according to P. Assume $R =< r_0, \cdots, r_{i-1}, r_i, r_{i+1} \cdots r_n >$ is an actual path used to route a packet from node r_0 to r_n at run time (Figure 3.5). R is obtained by sequentially looking up the routing tables from routers r_0 to r_{n-1}. Assume this path $R \notin P$ is a disallowed path which may cause circular dependencies with the other routing paths in P, then at least one edge $R_{\mu\nu}(\mu = r_{i-1}r_i,\ \nu = r_i r_{i+1})$ is an element of S_{edge} in Table-3.1 in order to generate the circular dependencies. However, since R is built from the routing tables in router r_i, which means there is an entry in the routing table of r_i using link $r_{i-1}r_i$ as input and link $r_i r_{i+1}$as output. It is already

known the entries of the routing tables are created according to the paths in the admissible path set P. Therefore, there is another path $p \in P$ while $p \neq R$ traversing through r_i from the input link $r_{i-1}r_i$ to the output link r_ir_{i+1}, hence p also contains the edge $R_{\mu\nu}(\mu = r_{i-1}r_i,\ \nu = r_ir_{i+1})$ in its CDG. On the other hand, since $R_{\mu\nu} \in S_{edge}$, it will create cycles in the CDG and P is actually not a deadlock-free path set. This contradicts with the assumption.

Now the issue is how to obtain the probability values $p(o)$ for each output o from the path traffic allocation ratios. For each router t and each of its input port d , we obtain the subset of admissible paths (*i.e.*,$T_{t,d}$) that pass through t using input port d. For a given destination tile b, using minimum path routing, only two candidate output ports ($o_1$and o_2) are feasible. The probabilities of selecting ports o_1 and o_2 are calculated by comparing the aggregate traffic of the paths inside $T_{t,d}$ that use ports o_1 and o_2, respectively, to route to the tile b. Using the notations in Table 3.1, let paths $l(s,b,k) \in T_{t,d}(o_1,s,b)$ and $l(s,b,l) \in T_{t,d}(o_2,s,b)$, we have:

$$p_{t,d}(o_1,s,b) + p_{t,d}(o_2,s,b) = 1 \tag{3.12}$$

$$\frac{p_{t,d}(o_1,s,b)}{p_{t,d}(o_2,s,b)} = \frac{\sum_{\forall l(s,b,k)\in T_{t,d}(o_1,s,b)} p(s,b) \times r(s,b,k)}{\sum_{\forall l(s,b,l)\in T_{t,d}(o_2,s,b)} p(s,b) \times r(s,b,l)} \tag{3.13}$$

After considering all the routers and the communication pairs, the port probability values are obtained and stored into the routing tables offline.

In order to maintain the global path traffic allocation ratios, each entry in the routing table $RT(t,d)$ needs to distinguish the source tile location s of the packet. Thus the number of entries in the routing table is increased. We can reduce the table size by grouping the entries of different sources but the same destinations together. The new format of the routing table becomes $RT(t,d) = \{(b,o,p(o)|o \in D_{out}(t), 0 \leq p(o) \leq 1\}$. In this case, the port probability value $p(o,b)$ in $RT(t,d))$ is calculated as:

$$p_{t,d}(o_1,b) + p_{t,d}(o_2,b) = 1 \tag{3.14}$$

$$\frac{p_{t,d}(o_1,b)}{p_{t,d}(o_2,b)} = \frac{\sum_s \sum_{\forall l(s,b,k)\in T_{t,d}(o_1,s,b)} p(s,b) \times r(s,b,k)}{\sum_s \sum_{\forall l(s,b,l)\in T_{t,d}(o_2,s,b)} p(s,b) \times r(s,b,l)} \tag{3.15}$$

Although the exact path traffic allocation ratios cannot be maintained, the traffic distribution on each router and each link can still be maintained which is more important for uniform energy distribution purpose. From the simulation results which will be presented in Section 6, this routing table implementation

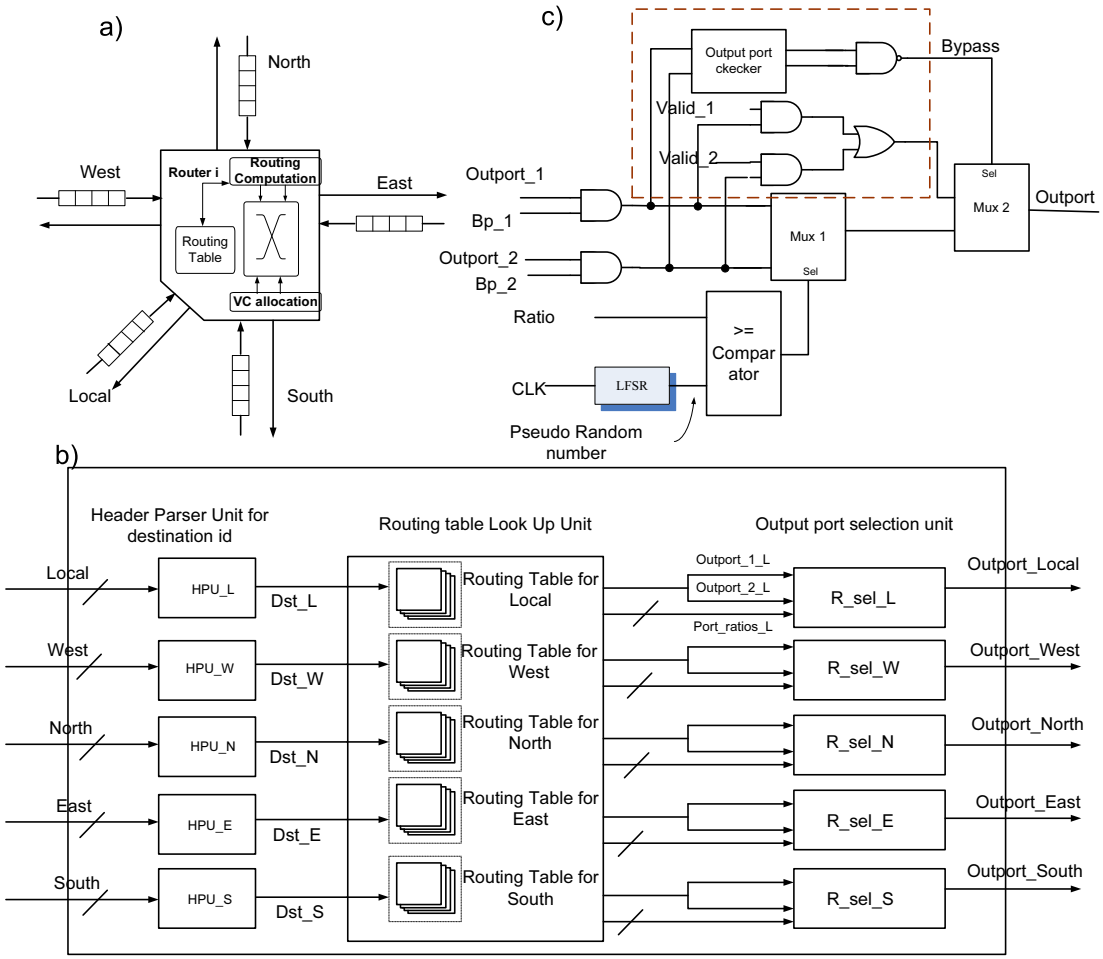

Figure 3.6: Block diagram of the router supporting ratio-based routing: a)Router microarchitecture, b)Routing computation unit, c)Ratio-based output port selection

achieves similar improvement in peak energy reduction and latency performance compared to the routing table using the source-destination pair.

3.5 Router microarchitecture

The block diagram of our proposed router to support thermal aware routing is illustrated in Figure 3.6-a. For minimum path routing, if the input direction and the destination are fixed, there are at most two candidate output ports, o_1 and o_2 , and $p(o_1) + p(o_2) = 1$. The routing selection unit in the router selects the output port O_{dir} by comparing the probability value $p(o_1)$ with a random number τ in $[0, 1]$. It will select port O_1 if $\tau \leq p(o_1)$, otherwise port O_2 will be chosen.

A pseudo random number generator using linear feedback shift register (LFSR) is employed to generate the random number. The header flits from each input

port are first decoded by a parser (the HPU module shown in Figure 3.6-b) to extract the destinations. Then by checking $RT(t, d)$, two candidate output ports are returned with the corresponding probability values. The output port selection unit then make a decision on the output port. In addition to the probability selection, backpressure information (Bp_1 and Bp_2 in Figure 3.6-c) from downstream routers are also used. If one candidate output port is not available due to limited buffer space, the backpressure signal will disable this output port from the selection.

3.6 Experimental results

3.6.1 Simulation environment setup

We implemented the proposed thermal-aware routing strategy for mesh-based NoC architecture. A C++ program is developed to analyze the traffic parameters automatically and generate the corresponding LP formulation for a given application. The LP problem is then solved by CVX [78] which is a toolbox embedded in MATLAB for convex optimization. In order to evaluate the NoC performance, a systemC based cycle-accurate simulator was developed combining the features of several widely adopted academic tools (including Noxim [79], Nirgam [80] and Booksim [61]). We used both synthetic traffic and real benchmarks to evaluate the performance and compare with other deadlock free routing algorithms (westfirst, northlast, negativefirst, oddeven and X-Y routing [6]). The real benchmarks include PIP (Picture-In-picture) [68], MWD (Multi-Window Display) [68], MPEG4 [81], VOPD (Video Object Plane Decoder) [68], MMS_1(Multimedia system mapping to 16 cores) [67], MMS_2(Multimedia system mapping to 25 cores) [82] and DVOPD (Dual Video Object Plane Decoder) [62]. For synthetic traffic, we use 8 different traffic scenarios, namely Uniform random, Transpose-1, Transpose-2 [83], Hotspot in center (Hotspot-C) [84], Hotspot in top-right corner (hotspot-Tr) [23], Hotspot in right-side (Hotspot-Rs) [23], Butterfly [79] and Bursty traffic [79]. 3×3, 4×4 and 5×5 NoC meshes are used for different benchmarks and 4×4 meshes are used for all synthetic traffics.

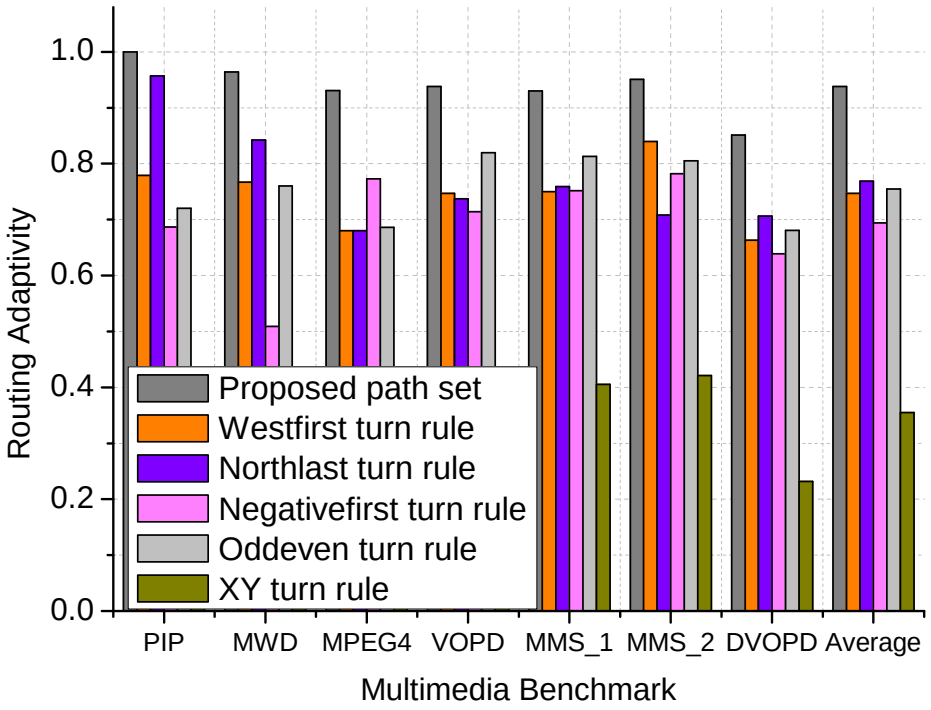

Figure 3.7: Adaptivity comparisons for different benchmarks

3.6.2 Adaptivity improvement

First we compare the routing adaptivity with other routing algorithm under several real benchmarks. The results are shown in Figure 3.7. It can be seen that 20%-30% more paths are available if we take the application traffic into consideration.

3.6.3 Peak energy simulation results

Next we evaluate the performance of the proposed algorithm in peak energy reduction. In the experiment, energy parameters ($E_{buffer-rw}$, E_{rc} etc.) are adopted from Noxim and Booksim simulator. After $10,000$ warm-up cycles, we carry out energy simulation using a fixed time window for all cases. Different simulations were done for each packet injection rate to obtain the average peak energy consumption. The simulation results for the real benchmarks and the synthetic benchmarks are shown in Figure 3.8 and Figure 3.9-3.11, respectively.

In the figures, the proposed source-destination routing table and destination-only routing table schemes are denoted as opt_source_dest and opt_dest, respectively. The theoretical optimal energy consumption which is obtained by solving

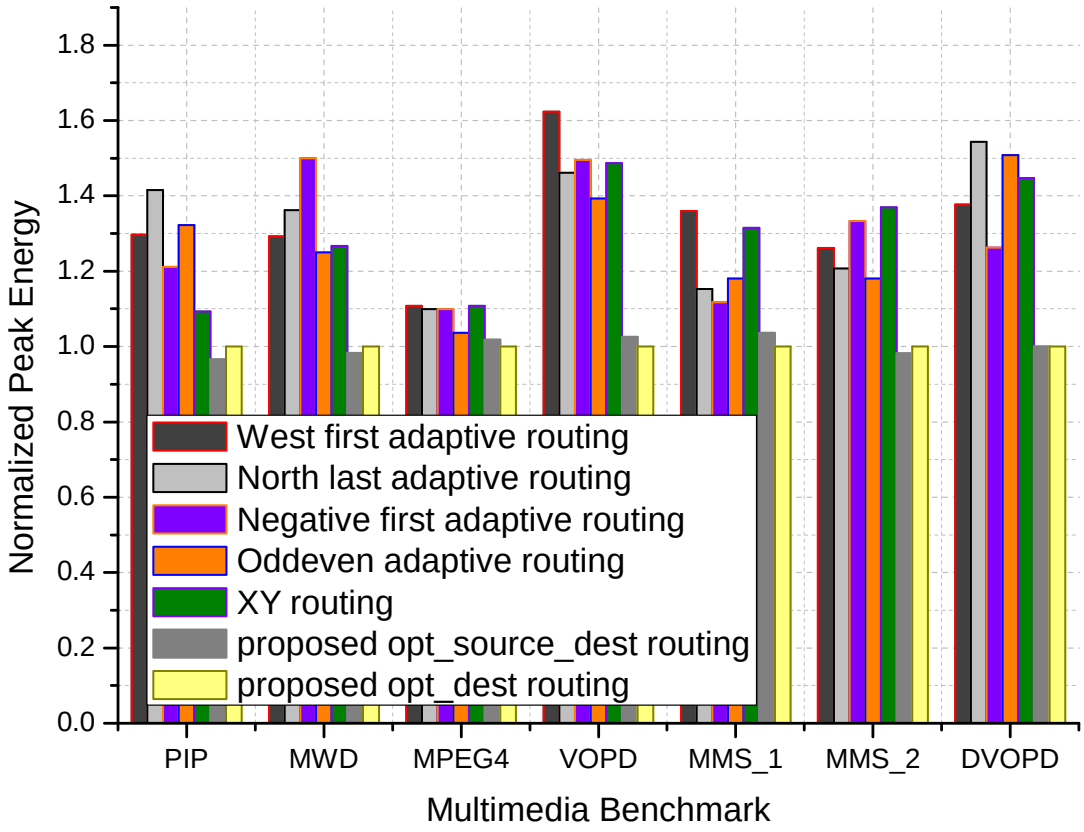

Figure 3.8: Peak energy simulation for real benchmarks

the LP formulation using CVX is also included as a reference and is denoted as opt_peak.

From Figures 3.8 and 3.9, we can see that comparing with other adaptive routing algorithms, the proposed routing algorithm can achieve more than 20% peak energy reduction in most traffics. At the same time, the latency performance can also be improved or maintained for all the cases.

In Figure 3.12, we illustrate a scenario of the tile energy distribution under Hotspot-C traffic (*i.e.,* four center cores are traffic hotspots). Figure 3.12-a is the energy distribution using the proposed thermal-aware routing algorithm, and Figures 3.12-b, -c and -d show the energy distribution using odd-even, negative-first and XY routing, respectively. For all the four cases, the total energy consumption of the NoC is same, *i.e.,* $4.55 \times 10^{-5} J$. However, the peak tile energy of the proposed, odd-even, negative-first and XY routing are $5.95 \times 10^{-6} J$, $8.09 \times 10^{-6} J$, $7.53 \times 10^{-6} J$ and $7.72 \times 10^{-6} J$ respectively. From the figures, we can see that our proposed scheme indeed leads to a more uniform energy profile across the NoC chip.

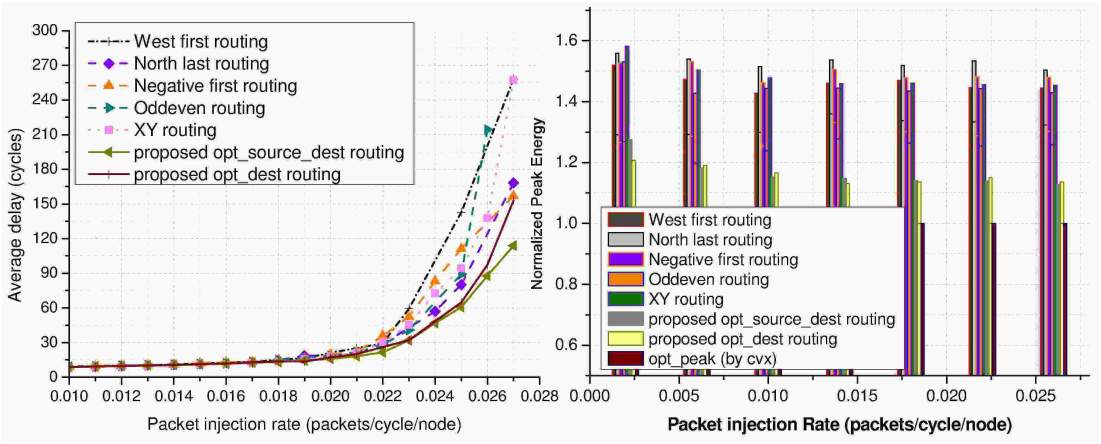

(a) Random traffic

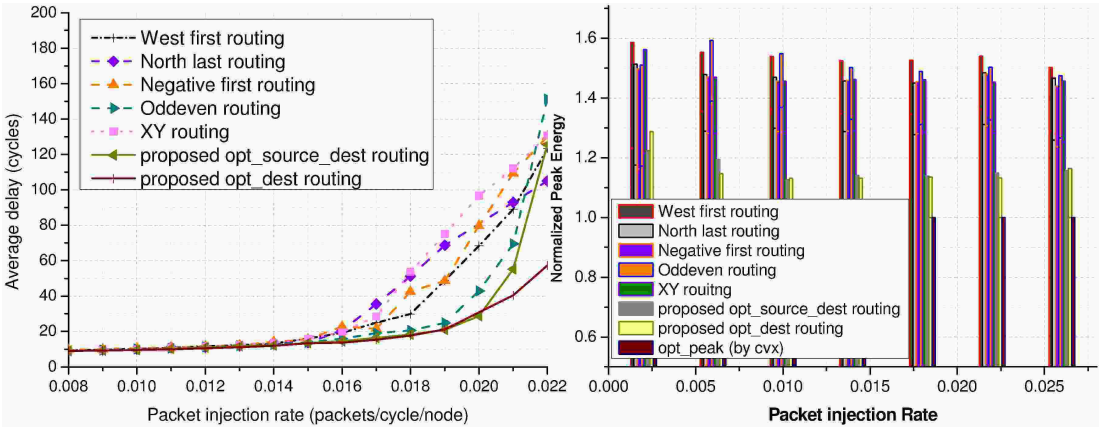

(b) Transpose-1 traffic

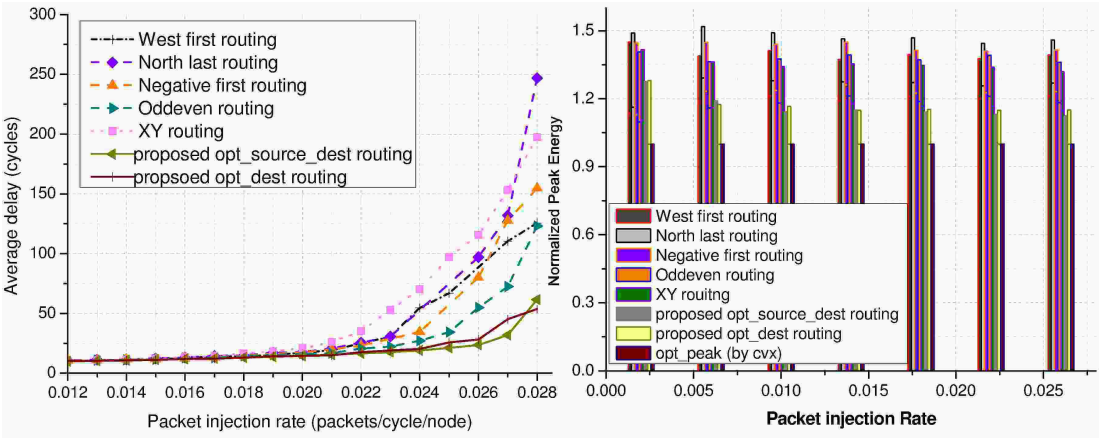

(c) Transpose-2 traffic

Figure 3.9: Latency and peak energy simulation for a) Random, b) Transpose-1 and c) Transpose-2 traffic

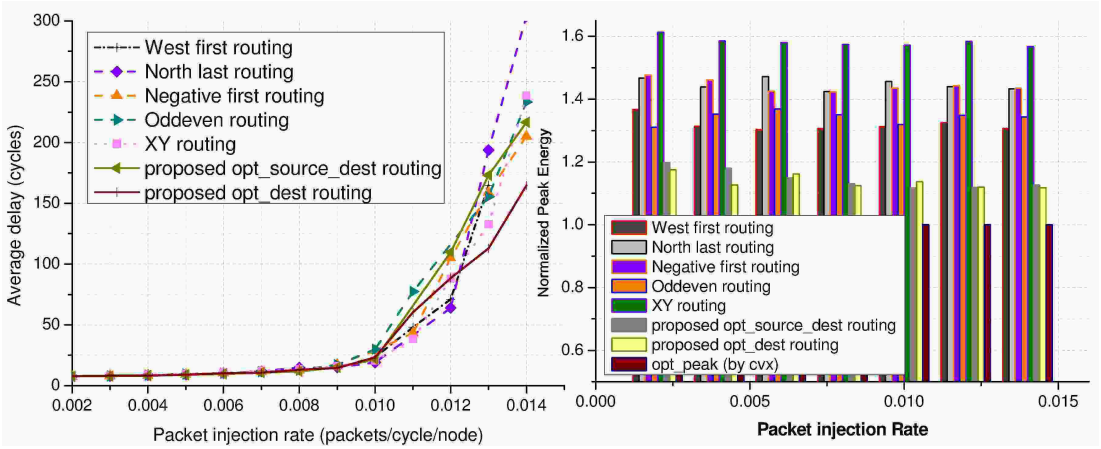

(a) Hotspot in right-side (Hotspot-Rs) traffic

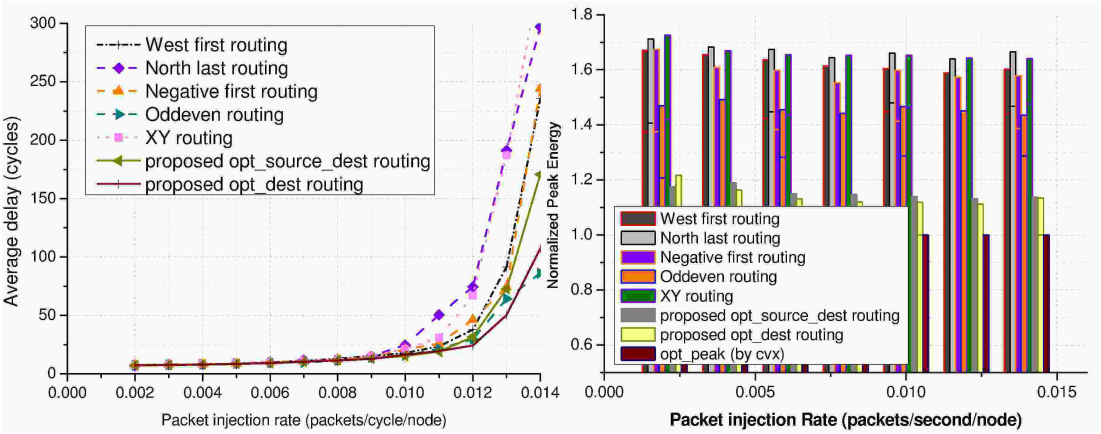

(b) Hotspot in top-right corner (Hotspot-Tr) traffic

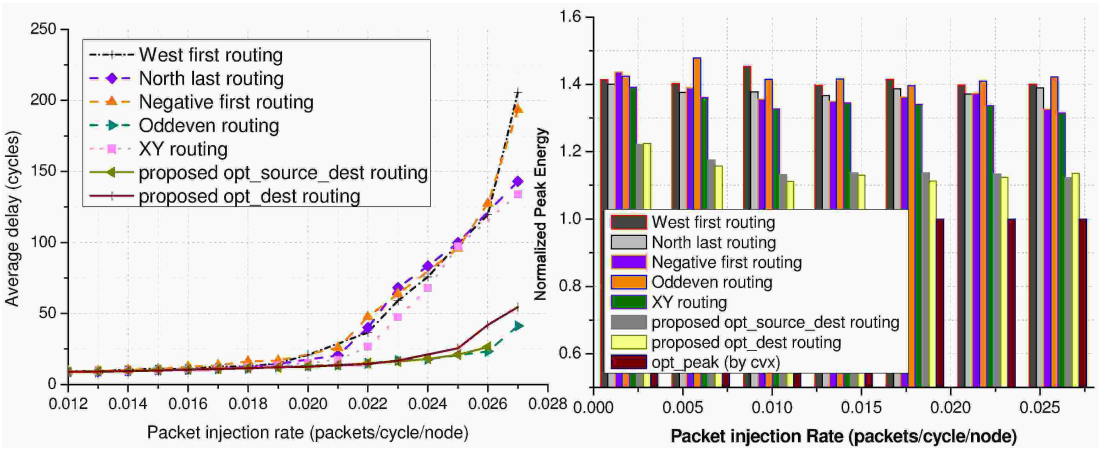

(c) Bursty traffic

Figure 3.10: Latency and peak energy simulation for a) Hotspot-Rs, b) Hotspot-Tr and c) Bursty traffic

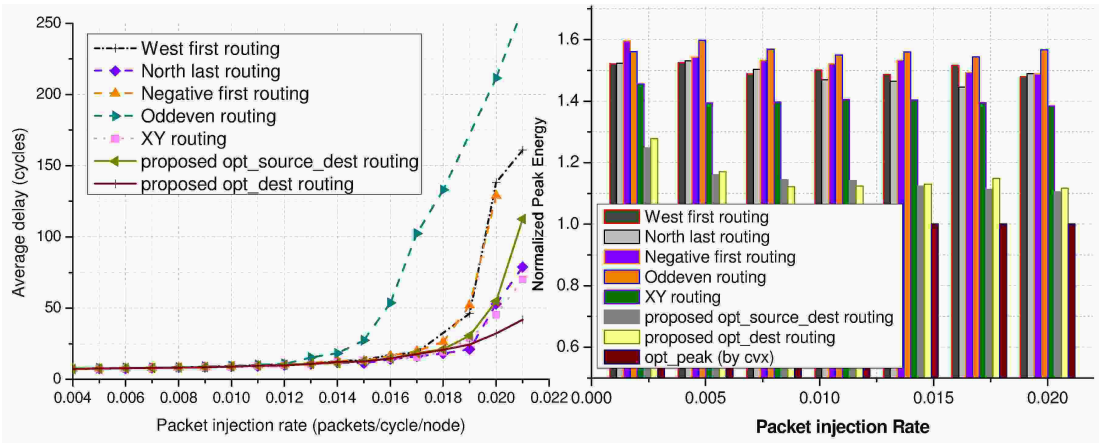

(a) Hotspot in center (Hotspot-C) traffic

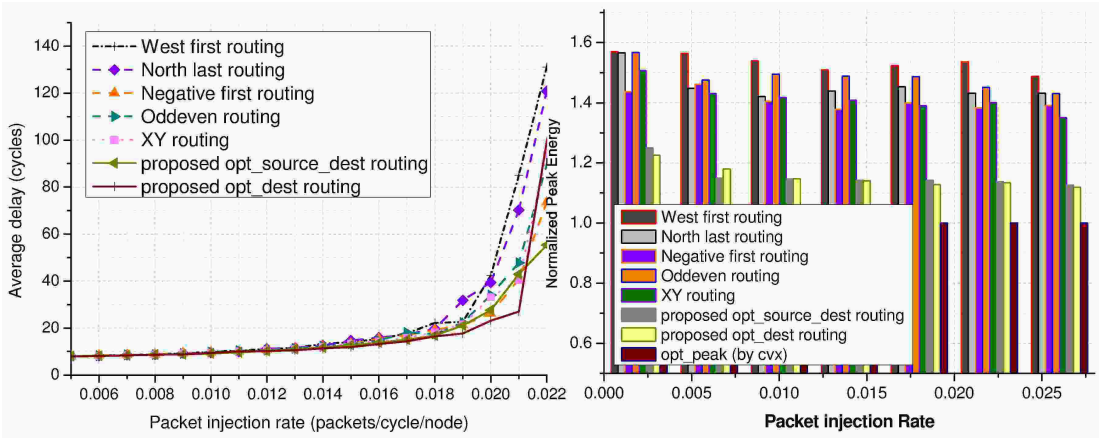

(b) Butterfly traffic

Figure 3.11: Latency and peak energy simulation for a) Hotspot-C and b) Butterfly traffic

In Table 3.3, we summarize the execution time of our off-line routing algorithm (including path generation and LP solving) under various mesh size and communication density ρ (ρ is defined as the ratio of the total number of communications pairs to the number of processors in mesh). It can be seen that the execution time is reasonable. For larger mesh size (7×7 or more) and more communication pairs (100 or more), the number of minimum paths increase dramatically. It takes longer time (2 -3 hours in average) to obtain the traffic allocation ratios. Since the traffic allocation is determined offline, in most cases, the time cost is still affordable. In case we want to reduce the execution time, we can restrict the number of minimum-length paths to a smaller subset.

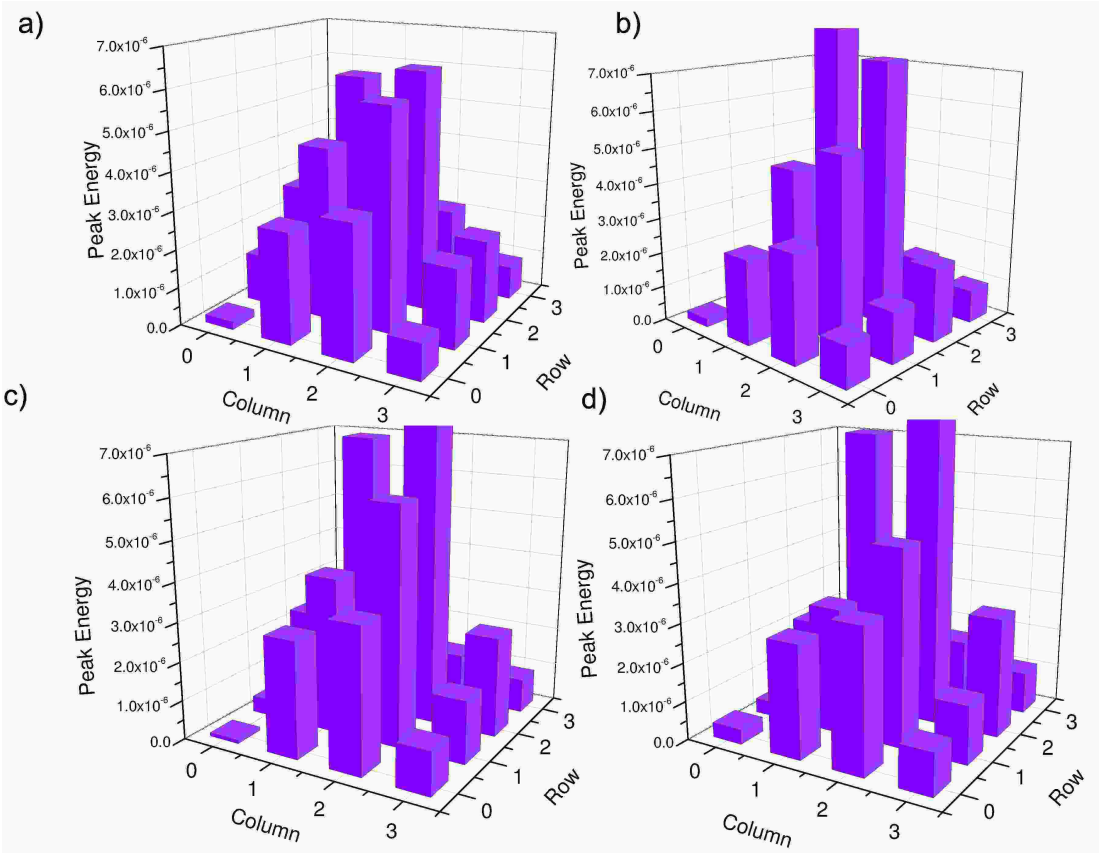

Figure 3.12: An example of the NoC energy profile under hotspot-4c traffic: a) proposed routing b) oddeven routing c) negativefirst routing d) XY routing

Table 3.3: Execution time for different mesh sizes

	Running time (s)					
Mesh Size	3×3		4×4		5×5	
Com density (ρ)	2	4	2	4	2	4
Time (s)	5	6	18	20	227	243

3.6.4 Simulation results with different processor/router energy ratio

The energy consumption of a tile consists of that of the processor and the router. For different applications, the power contribution of the router and the processor varies greatly. Our thermal-aware routing algorithm can only re-distribute the router power. We want to evaluate the effectiveness of our algorithm on the peak energy reduction when the ratio of the energy contribution from the router and the processor varies. Let $r_e = \frac{Average\ processor\ energy}{Average\ router\ energy}$. The experimental results shown in the previous sections assume $r_e = 1$. In this sub-section, we simulate

Table 3.4: Peak energy reduction under various r_e(for Random, Transpose-1, Transpose-2 and Hotspot-C traffic)

Benchmark	Peak energy reduction							
	Random		Transpose-1		Transpose-2		Hotspot-C	
Average Energy Ratio (r_e)	vs. XY	vs. OE	vs. XY	vs. OE	vs. XY	vs. OE	vs. XY	vs. OE
0.67	17.4%	12.8%	15.6%	17.9%	9.5%	12.1%	15.7%	17.6%
1.00	15.7%	15.3%	13.3%	14.2%	10.3%	10.3%	14.4%	16.2%
1.33	13.3%	9.5%	12.5%	14.5%	7.2%	9.4%	13.2%	14.9%
1.67	10.6%	8.6%	11.2%	13.8%	6.4%	9.4%	12.3%	14.0%
2.00	11.9%	9.3%	11.1%	13.8%	6.5%	9.1%	11.6%	15.0%
2.33	10.5%	8.2%	8.5%	10.7%	5.5%	7.0%	10.8%	13.9%
2.67	9.4%	7.7%	8.9%	11.1%	4.3%	6.9%	10.2%	13.0%
3.00	8.8%	7.0%	8.8%	10.6%	3.7%	5.6%	9.6%	10.9%
3.33	7.9%	5.9%	7.1%	9.9%	4.2%	6.6%	9.1%	10.5%
3.67	8.1%	6.5%	7.8%	9.2%	4.1%	5.8%	8.7%	9.9%
4.00	7.9%	6.1%	7.0%	9.0%	4.6%	6.4%	8.3%	9.4%
Average	11.04%	8.81%	10.16%	12.25%	6.03%	8.05%	11.26%	13.21%

the peak energy reduction using our routing scheme for different r_e values. Tables 3.4, 3.5 and 3.6 summarize the results for synthetic traffic and real benchmarks, respectively. From the results we can see that when r_e increases, the peak energy reduction is smaller as the relative contribution of the router energy reduces. Overall, we can achieve an average 6%-17% peak energy reduction over other two existing routing schemes when r_e ranges from 0.67 to 4 across all the benchmarks.

3.7 Conclusion

NoC has been widely adopted to handle the complicate communications for future MPSoCs. As temperature becomes one key constraint in NoC, in this chapter, we propose an application-specific and thermal-aware routing algorithm to distribute the traffic more uniformly across the chip. A deadlock free path set finding

Table 3.5: Peak energy reduction under various r_e(for Butterfly, Hotspot-Tr,Hotspot-Rs and Bursty traffic)

Benchmark	Peak energy reduction							
	Butterfly		Hotspot-Tr		Hotspot-Rs		Bursty	
Average Energy Ratio (r_e)	vs. XY	vs. OE	vs. XY	vs. OE	vs. XY	vs. OE	vs. XY	vs. OE
0.67	14.7%	20.2%	16.6%	9.9%	21.4%	11.2%	12.6%	16.1%
1.00	13.1%	19.1%	15.4%	9.2%	19.9%	10.5%	10.4%	13.0%
1.33	10.8%	15.2%	13.6%	7.9%	18.5%	9.7%	8.1%	11.3%
1.67	9.8%	14.4%	12.5%	7.3%	17.2%	8.9%	7.5%	9.9%
2.00	9.9%	13.7%	11.6%	6.7%	16.3%	9.4%	5.6%	9.1%
2.33	8.9%	13.0%	11.0%	6.5%	15.5%	8.2%	7.1%	8.7%
2.67	8.8%	12.8%	10.1%	5.8%	14.6%	7.8%	5.3%	7.9%
3.00	8.7%	11.2%	9.5%	5.4%	13.8%	7.2%	5.5%	7.6%
3.33	7.4%	9.8%	8.9%	5.1%	13.1%	6.9%	4.9%	7.2%
3.67	5.8%	9.4%	8.4%	4.8%	12.5%	6.6%	4.3%	6.0%
4.00	6.0%	8.8%	8.0%	4.6%	12.0%	6.3%	4.5%	6.4%
Average	9.45%	13.42%	11.42%	6.65%	15.89%	8.43%	6.89%	9.38%

Table 3.6: Peak energy reduction under various r_e (for real benchmark traffic)

Benchmark	Peak energy reduction					
	MMS-1		MPEG4		VOPD	
Average Energy Ratio (r_e)	vs. XY	vs. OE	vs. XY	vs. OE	vs. XY	vs. OE
0.67	17.7%	11.8%	7.8%	10.2%	16.5%	28.9%
1.00	15.7%	10.4%	6.9%	9.0%	12.3%	23.6%
1.33	14.5%	9.8%	6.1%	8.0%	10.3%	20.5%
1.67	12.5%	8.1%	5.5%	7.3%	9.9%	19.0%
2.00	10.6%	6.5%	5.4%	7.1%	9.3%	17.7%
2.33	9.3%	5.6%	5.0%	6.5%	8.3%	16.1%
2.67	10.4%	7.0%	4.6%	6.0%	7.6%	14.8%
3.00	9.7%	6.5%	4.3%	5.6%	7.6%	14.2%
3.33	8.5%	5.6%	4.0%	5.5%	7.7%	13.8%
3.67	7.9%	5.2%	4.2%	5.4%	7.1%	13.0%
4.00	7.3%	4.7%	3.9%	5.0%	6.5%	12.0%
Average	11.28%	7.38%	5.25%	6.87%	9.37%	17.6%

algorithm is first utilized to maximize the routing adaptivity. A linear programming (LP) problem is formulated to allocate traffic properly among the paths. A table-based router is also designed to select output ports according to the traffic allocation ratios. From the simulation results, the peak energy reduction can be as high as 20% for both synthetic traffic and real benchmarks.

CHAPTER 4

FAULT-TOLERANT ROUTING SCHEME FOR NOCS USING DYNAMIC RECONFIGURATION OF PARTIAL-FAULTY ROUTING RESOURCES

In this chapter, we propose a fault-tolerant framework for Network-on-Chips (NoC) to achieve maximum performance under fault. A fine-grained fault model is first introduced. Different from the traditional link or node NoC fault models which assume the faulty resource is totally unfunctional, we distinguish the faulty components and handle them according to their fault classes. By doing so, we can avoid unnecessary partitioning of the network and hence achieve a higher connectivity under high fault rate. In particular, two new dynamic re-configuration schemes at the router level, namely Dynamic Buffer Swapping (DBS) and Dynamic MUX Swapping (DMS), are proposed to deal with the buffer and crossbar faults, which are the main sources of failure in the router. In these schemes, the healthy resources in the router are maximally utilized to mitigate the faults. A deadlock recovery scheme is designed to handle the potential deadlock hazard due to the proposed DBS operation. Experimental results show that we can achieve higher packet acceptance rate as well as lower latency compared with state-of-the-art fault-tolerant routing schemes

4.1 Introduction

For NoC-based Multiprocessor-System-on-Chip (MPSOC), the complex system is highly susceptible to the prominence of faults [86]. A significant error tolerance to both permanent and temporary faults is thus strongly desired. Temporary errors are mainly due to the on-chip crosstalks and coupling noise while permanent faults are mostly caused by the worn-out devices, manufacturing defects and accelerated aging effects [10]. The NoC architecture offers a good opportunity of achieving high reliability since it provides multiple paths between the source-destination communication pairs and allows reconfiguration under faulty conditions.

To achieve fault tolerance in NoC, a proper fault model is required for the clear identification and isolation of the defective components so that the routing algorithm can maximally utilize the remaining functionality to achieve the optimal performance. Traditional NoC fault model treats the router fault as a node fault [85] or a link fault [9,19,86] and assumes total outage of the faulty resource. However, a fine-grained functional fault model will be more beneficial because graceful degradation of network operation by exploiting the partially functional resources can be achieved [10]. In this chapter, we identify the major fault components in the NoC and propose a fine-grained permanent fault model, which distinguishes the partial router faults (buffer or crossbar faults) from the hard link faults. Hard link faults may lead to a partitioned network. On the other hand, a faulty router with an appropriate reconfiguration strategy, may be able to tolerate the buffer or crossbar errors and still supports the full connectivity in the NoC with reduced capacity. This motivates a need to develop an integrated fault handling methodology for different fault classes. An example comparing different fault models given the same fault pattern is shown in Fig. 4.1. As shown in the figure, if we adopt the generic node or link fault model (Fig. 4.1-a and -b) which do not consider whether the fault is caused by the buffer or the crossbar, we may isolate many nodes from the network. To increase the connectivity under faulty situation, [19] adopts a static offline port swap algorithm to reduce the number of faulty links (replacing port P_a with port P_b in this example). However, the resultant network is still partitioned, because the port fault is treated as a hard link fault under this algorithm.

In this chapter, we propose a dynamic swapping scheme which shares the healthy resource in the router for different ports at run-time. As shown in Fig. 4.1-d, by leveraging the other healthy buffers or multiplexers (Muxes) and dynamically re-allocating the resources, the network connectivity can remain intact which result in better performance in terms of packet acceptance rate. Two new techniques (DBS and DMS) are proposed to handle the buffer and crossbar faults, respectively. An approach that combines the turn model based routing and an efficient deadlock recovery scheme is also proposed to resolve the potential deadlock issue when the routers are working in the DBS mode.

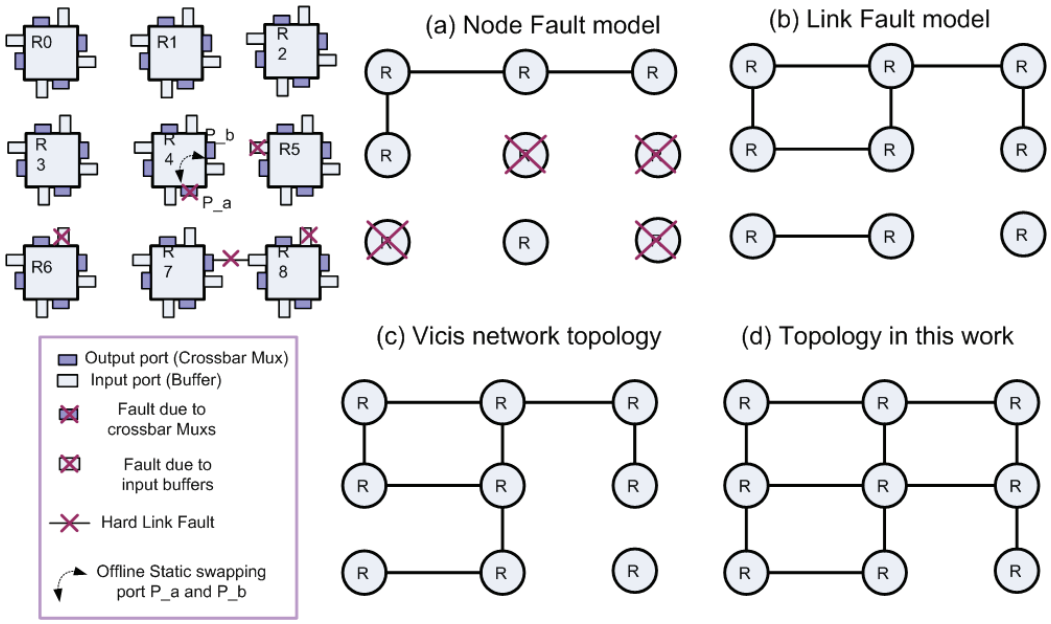

Figure 4.1: A Comparison of various fault model

4.2 Background

The growing concern on reliable NoC has prompted extensive researches in this area, from self-detect and self-diagnosis routers [9] to fault-tolerant routing algorithm design [86]. In [9], a router with built-in-self-test (BIST) and self-diagnose circuits is proposed to detect and locate the faults in the FIFOs and the crossbar MUXes. After identifying the faults, the FIFO faults and the MUX faults are treated as input and output link faults, respectively. This fault categorization method reduces the connectivity of the network and hence degrades the performance if there are a large number of faults. In [10], a detailed online fault diagnosis framework is proposed to diagnose the link, crossbar and arbitration/allocation logic faults. The remaining healthy paths are identified from the faulty paths and the routing algorithm utilizes these healthy paths to achieve graceful degradation. However, the faulty paths can still lead to network partitioning as they are determined statically. Moreover, this work is designed for buffer-less router and may not be suitable for buffered router architecture, which is required for high throughput and low latency design.

There have been many fault-tolerant routing algorithms proposed to deal with faulty components in NoC. In [85], the faulty router is treated as a node fault and the routing algorithm can be dynamically reconfigured to re-route the packets following the contour around the faulty nodes. This region-based routing algorithm

is only suitable for one-faulty-router topology. Also, the node fault model may cause the faulty router totally isolated from the network, while in real situation, the faults may only affect a certain part of the router and the rest of it can still function correctly. In [86], the NoC faults are modeled as link level hard failures and a highly resilient routing algorithm is designed to reconfigure the routing table in an offline process. The algorithm is static in nature and there exists a trade-off between the network connectivity and the deadlock avoidance by selectively removing the turn rules. In [87], a deflection routing algorithm is proposed to dynamically re-route the packets towards the destinations based on the stress factors which reflect the traffic condition. The NoC faults are treated as the link faults and it does not clearly describe how to avoid deadlock and livelock based on the stress factors.

4.3 Fault-tolerant NoC design

In the context of NoC, the most common source of failures includes link errors and single-event upsets within the router [88]. Since the input buffers and the crossbar occupy the majority of the router area (67.68% for the FIFO buffer and 11.14% for the crossbar MUX in [9]) , most of the faults in the router may occur there [19]. In this chapter, first we propose a hybrid fine-grained fault model to distinguish the nature and consequence of different faults and briefly describe the fault-diagnosis technique to distinguish the types of fault [89]. Then, we introduce the DBS and DMS techniques which dynamically reconfigure the routers to handle buffer faults and crossbar faults while still maintaining the maximum service of the router. After that, the scheme to deal with link faults is discussed. Finally, we introduce the deadlock recovery mechanism used in our framework.

4.3.1 Fault-diagnosis and hybrid fault model

Fault-diagnosis infrastructure

The fault diagnosis techniques not only should detect the occurrence of faults but also should indicate precisely the type as well as the position of the faults. In this work, we assume the NoC fault diagnostic technique proposed in [9, 10] is used

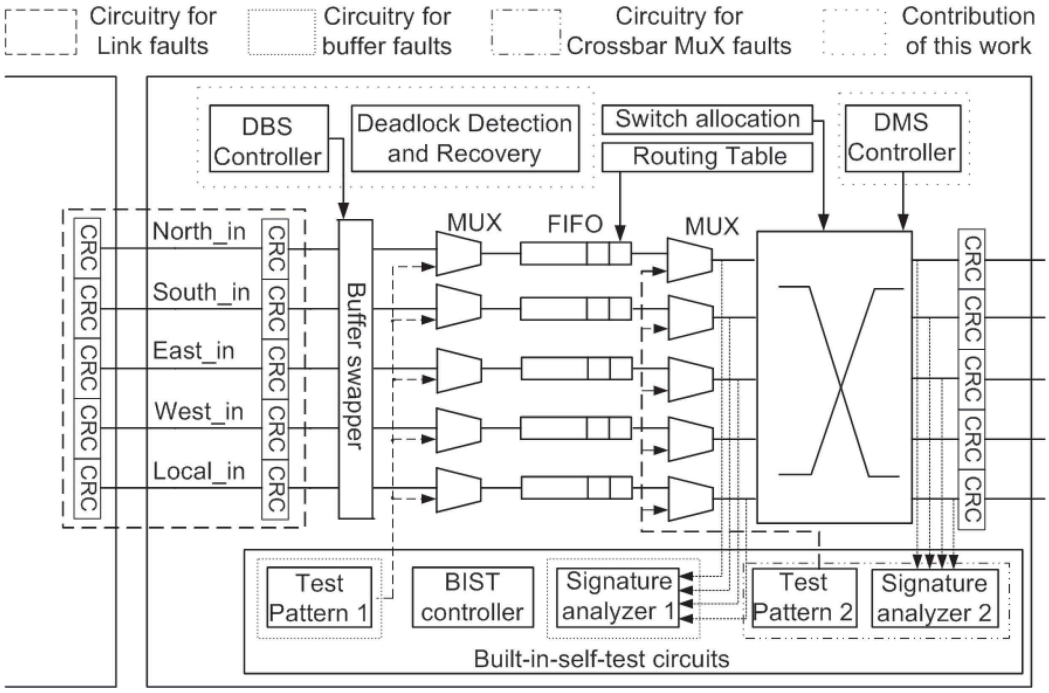

Figure 4.2: Architecture of fault-tolerant NoC [9, 10]

(shown in Fig. 4.2), which provides the detail information of the faults in router. In the following, we briefly review the state-of-the-art fault-diagnosis structures in NoC [9, 10, 19].

As shown in Fig. 4.2, the cyclic redundancy check (CRC) units are used for the link fault diagnosis. The packet correctness is verified by the CRC modules associated with each input port. The link transmission errors is characterized by a corresponding checksum mismatch [10]. Once a switch-to-switch transmission error is detected, the BIST controller in the upstream router will start the link diagnosis by sending the test patterns to the downstream neighbor and find out the specific link wires that have faults. Similarly, the intra-router error will be discovered if the packets pass the CRC check at the input of the router but fail at the output side [10]. On the detection of the internal fault, the router will switch to the test mode to diagnose the specific faulty components such as the input buffers or switch crossbars. As illustrated in Fig 4.2, the Test Pattern Generator (TPG) will generate the test vectors for the buffers (TPG1) and crossbar MUX (TPG2) respectively, while the Signature Analyzer will compare the results and identify the specific faulty components in either FIFO buffers or the crossbar. The diagnosis results are finally forwarded to the DBS and DMS units for router reconfiguration based on the proposed hybrid fault model.

Table 4.1: The proposed hybrid model for permanent faults

Category	Component	Handling scheme	Impacts
Static	Link	fault-tolerant routing to avoid faulty links	Partitioned network
Dynamic	Buffer	time-multiplexed reuse other buffers	Reduced bandwidth / deadlock
	Crossbar	time-multiplexed reuse other MUXs	Reduced bandwidth

A novel hybrid fault model

Table 4.1 summarizes our proposed hybrid fault model. The permanent faults are divided into two classes: the static and the dynamic fault. Faults occurred at the interconnect wires is denoted as link fault and is classified as a static fault, which physically disconnects two nodes. Of course in real situation, not all the wires are faulty and we can still exploit the healthy wires to maintain the connectivity similar to the approaches proposed in [90]. In this chapter, we mainly focus on the recovery of the router faults and hence we assume a totally-failure link fault. To handle the static fault, a fault-aware routing algorithm is exploited to bypass the error links. On the other hand, if the fault occurs within the buffer or the crossbar, we can mitigate them by reusing the other healthy resources in the router. All the datapath will remain intact with a small penalty of reduced bandwidth. This kind of fault is categorized as dynamic faults in our model.

4.3.2 Dynamic buffer swapping (DBS)

Dynamic buffer swapping (DBS) algorithm is proposed to handle the input buffer faults. For traditional fault model, the buffer fault is treated as a hard link fault and data cannot be transmitted even the link is healthy. In our approach, through proper swapping of another substitute buffer to hold the packets destined for the faulty buffer, a graceful performance degradation can be achieved by fully utilizing the link. By adopting DBS algorithm, no extra buffers or virtual channels are needed in the router, which satisfies the stringent area constraint of NoC [85].

DBS operation

DBS technique operates with three states. As shown in Figure 4.3, assume fault occurs at the north input buffer of router R1, packet 0 is blocked in router R2. We denote this situation as state S0. After receiving the transmission request token, R1 enters the port-swap scheduling state S1. In S1, R1 will arrange another input port buffer to substitute the north buffer based on a predefined swapping strategy. Assume the east input buffer is chosen in this example. Then a traffic control token is propagated to router R4 indicating that the transmission link (R1-R4) will be shut down in the next state for DBS operation. Upon the reception of this flow control token, R4 will stop allocating and granting new requests to traverse through its west output port after the tail flit of packet 2 finishes routing. R1 will also continuously monitor its east input buffer. Once the buffer serves the tail flit of packet 2, it will enter the buffer swapping state S2. In this state, the communication between R2-R1 link is enabled while link R4-R1 is currently disabled. The packet 0 can thus be transmitted to R1's east input buffer and further be forwarded towards its destinations. After packet 0 departs R1, R1 returns to normal state S0 and link R4-R1 is re-enabled again. The router will keep running at state S0 until another transmission request is asserted to use R1's north input buffer.

Avoid packet interleaving in DBS

One situation needs to be addressed and avoided is the packet-level interleaving when DBS operation is performed. In the previous example, in S1, R4 will control its traffic towards its west output port temporarily. Without this control, if R4 continues its transmission for the following packets (packet 3 in Fig. 4.4 for instance) and R1 jumps to state S2 for buffer swapping operation, the east buffer of R1 will store part of the flits from packet 3 and the whole packet 0 from R2. The remaining flits of packet 3 are still in R4 waiting to be routed. Thus packet 0 will interleave with packet 3 in the east input buffer of R1. The header of packet 0 will erase the output reservation information created for packet 3. After packet 0 is sent and R1 returns to the initial state S0, the residual flits of packet 3 will enter R1 from R4. These flits lose the routing information in R1 which is previously reserved by their header and have to be discarded.

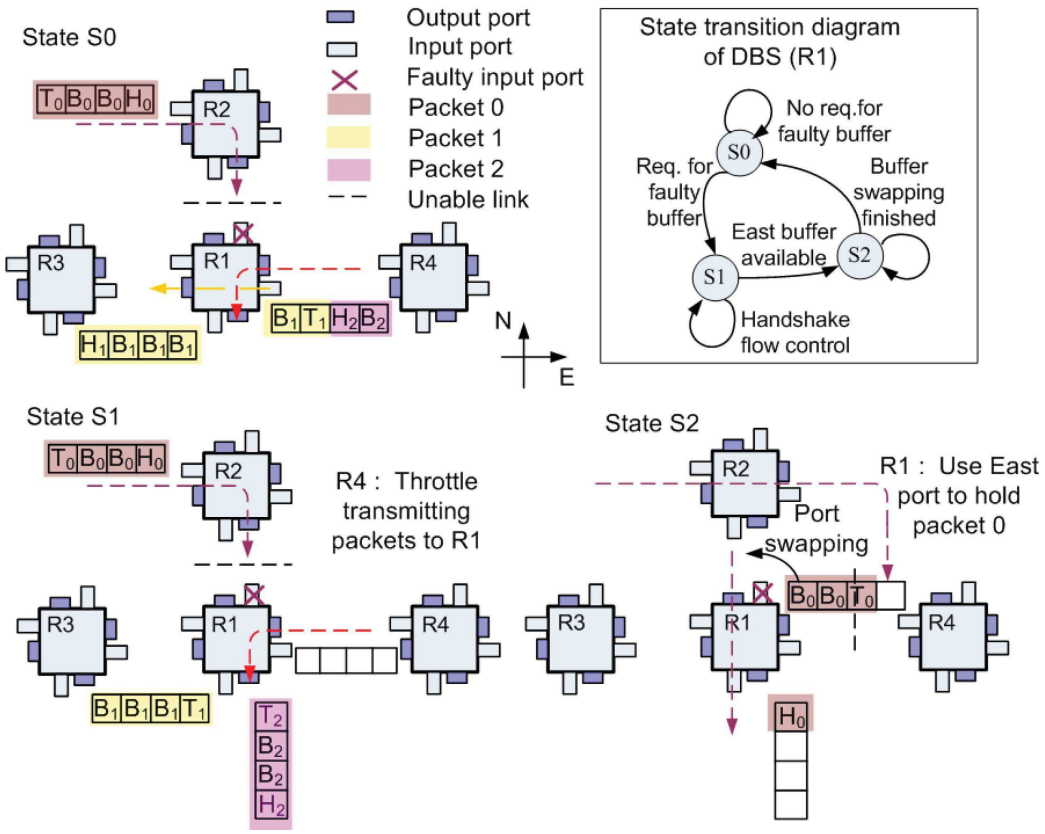

Figure 4.3: Dynamic Buffer Swapping (DBS) technique

Figure 4.4 illustrates the above situation. In order to avoid interleaved packets due to the DBS operation, a simple but effective flow control scheme is developed. Two handshake signals, current transmit status (CTS) and current port status (CPS), are associated with each port. Both CTS and CPS signals are sent by the faulty routers and read by their neighbors. The CTS (1bit) token manages the link transmission. The router can send the packets across the link only when it receives a '1' in the CTS. The CPS (2 bits) token is responsible for intra-router transmission. If "00" is read from one port, the port can be used as an normal output port for all the packets. When "01" is read, the router (R4 in Fig. 4.3) will begin monitor this port and grant it only to the packets which have been partially transmitted across the link (packet 1 and 2 in Fig. 4.3) to avoid the packet-level interleaving. Once the residual tail flit (T_2 in Fig. 4.3) is received, the faulty router enters buffer swapping state, and sends CPS "10" to its swapped port (east output port of R1 in Fig. 4.3) and the corresponding neighbor (R4 in Fig. 4.3) will not transmit packets to this link during the whole state. Within each router, three registers are used to record the faulty port (FP_R), the scheduled swapping port (SP_R) and the disabled port (DP_R), respectively. The status of

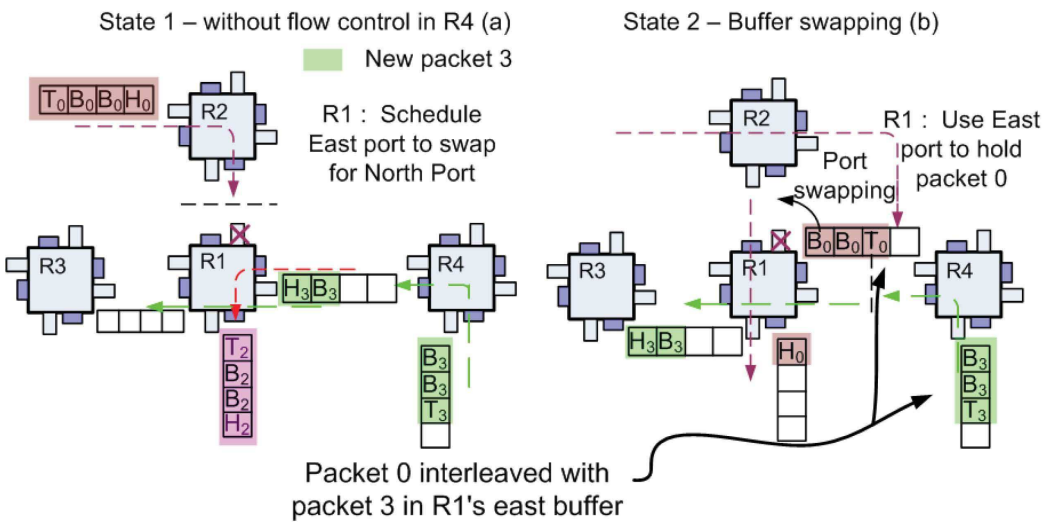
State 1 – without flow control in R4 (a)
New packet 3
R1 : Schedule East port to swap for North Port
State 2 – Buffer swapping (b)
R1 : Use East port to hold packet 0
Port swapping
Packet 0 interleaved with packet 3 in R1's east buffer
R1
R2
R3
R4

Figure 4.4: A pathological case for packets interleaving

these tokens for the example shown in Fig. 4.3 are summarized in Table 4.2.

buffer swapping strategy design

In DBS we need to schedule a healthy port to swap with the faulty port. As shown in Fig. 4.3, a buffer swapper is needed to match the physical link connections with the swapped healthy FIFO input. We can use a five-by-five port swapper to connect each physical link with every input buffer, similar to the swapper used in [19]. This architecture provides the highest flexibility for buffer swap scheduling. A round-robin DBS scheduling can be adopted by dynamically using different buffers to swap with the faulty port so to even out the bandwidth penalty for each input. For example, if the north port is faulty and in the previous swapping state, we use the east port to hold the packets from the north, then we will choose to use the south port to handle the next request for the north port.

A more area-efficient implementation for buffer swapping is to use a dedicated swapping strategy. This will reduce the area overhead of the port swapper as we do not need to have a crossbar-like connection architecture. The faulty port can only choose its predefined neighbor for swapping operation. From the simulation results, which will be presented in the next section, it is shown that this swapping strategy achieves only a little performance penalty compared with the round-robin DBS.

Table 4.2: Handshake flow control for DBS

States	Signals status (North port of R1 is faulty)	Operations / representation on tokens
S0	R1 North writes: CPS "10"; CTS "0"	R2 blocks packet 0
	R1 Reg: FP_R: "N"; SP_R: "ϕ";DP_R: "N"	R1 north port is faulty
S1	R1 North writes: CPS "10"; CTS "0"	R2 blocks packet 0
	R1 East writes: CPS "01";CTS "1"	R4 monitors request to the west
	R1 Reg: FP_R "N"; SP_R: "E"; DP_R: "N"	R1 prepares east to substitute north
S2	R1 North writes: CPS "00"; CTS "1"	R2 transmits packet 0 to east
	R1 East writes: CPS "10";CTS "0"	R4 blocks transmission to R1
	R1 Reg: FP_R "N"; SP_R: "E"; DP_R: "E"	R1 east port buffer is used for north link

4.3.3 Dynamic MUX swapping (DMS)

The crossbar in the router is usually implemented in the form of five 4×1 MUXs [9]. As shown in Fig. 4.5, if one of the MUX malfunctions, the data will not be presented to the output link correctly. Here we present a dynamic MUX swapping (DMS) technique to mitigate the MUX faults and make all the output links available for routing. The proposed ring-based crossbar swapper is shown in Fig. 4.5. The ring topology allows each MUX to be shared among its neighboring output ports. If the MUX corresponding to the east output port is faulty, the MUX for the north port can be time shared with the north and east ports by turning on-off the switches in MUX swapper accordingly. By doing so, all the output links can be maintained.

4.3.4 Handling hard link faults

For the link wire faults, since the physical connectivity between the routers has been damaged, we need to use a robust routing algorithm to bypass the faulty links. In this work, we use a similar approach as proposed in [19, 86] to route the packets under the existence of hard link faults.

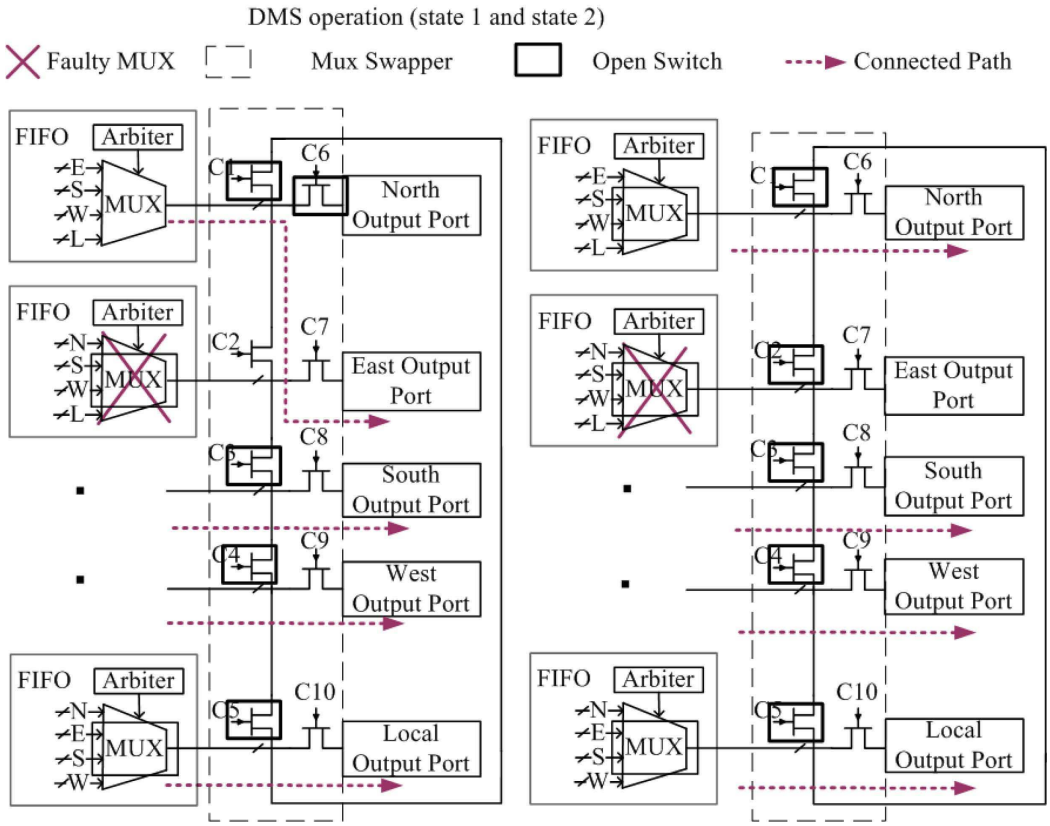

Figure 4.5: DMS hardware implementation

4.3.5 Deadlock issue and handling strategy

Deadlock in proposed framework

To ensure deadlock free routing, in normal operation, Oddeven (OE) routing algorithm [91] is used in our framework. However when faults occur, due to the DBS or the re-routing, a deadlock may be introduced.

Figure 4.6 gives an example of two deadlock situations. In Fig. 4.6-a, a local deadlock will occur between the two nodes R1 and R2 if the south port of R2 and north port of R1 are faulty and all the packets in other ports choose these two ports as outputs. In this case, R1 (R2) will be stuck at state S1 waiting for the original packets in the substitute port of R2 (R1) being completely transmitted. In Fig. 4.6-b, a global blockage situation is illustrated assuming the south to east turn is forbidden. If the scheduler in R2 arranges its west input buffer to hold packets from north and the scheduler in R1 arranges its east input buffer to substitute the faulty south buffer, a circular dependency on buffer resources can thus be generated as in the figure. According to [92], this kind of circular

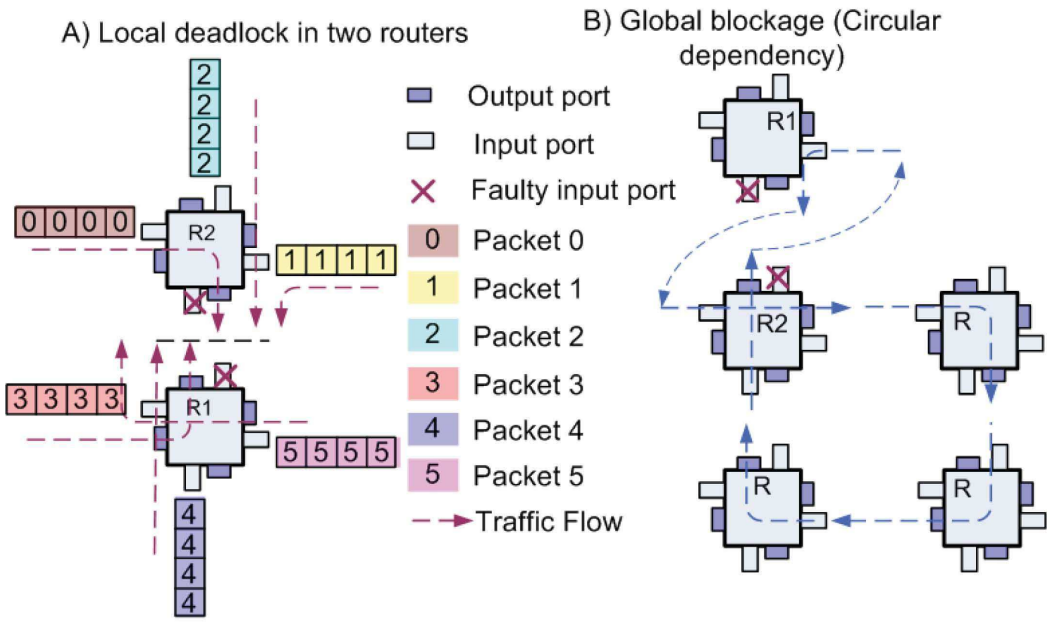

Figure 4.6: Deadlock situation for DBS operation

dependency on the NoC resources may bring deadlocks and stall the whole system.

For DMS operation, we only time-share the MUX between two output ports and do not change the dependency of the buffers in the routing algorithm. Hence it will not create deadlock if the original routing algorithm is deadlock-free.

Strategies to relieve deadlock

We employ a detection scheme, similar to the DISHA approach in computer network [93] to detect the blockage in the network. When the router receives the header flit of a packet, an embedded counter is triggered to count the number of cycles it stays in the buffer. The counter continues to increment until the tail flit has been successfully transmitted or its value exceeds a predefined threshold T_{sh}. In the latter case, the router will enter the deadlock handling mode. After that the counter would be reset and waits for next packet.

In this chapter, we propose a set of strategies to relieve the deadlock situation. We mainly deal with the packets in the faulty routers in DBS mode because DMS and the routing algorithm will not generate circular dependency. First, we set the SP_R register to the next port so as to initialize another swapping operation using a different substitute port. In Fig. 4.6-b, by using the east input buffer of R2 to swap with the north faulty buffer, we can break the original dependency and avoid the blockage. If all the ports in the router are not available for swapping

as the case in Fig. 4.6-a, we will reset the routing computation (RC), switch allocation (SA) information of these packets to request a different output port. In figure 4.6-a, if packet 1 chooses the west instead of south as output direction and leaves R2 successfully, then the deadlock problem can be solved accordingly.

If the re-computation step fails to relieve the blockage, and the local processing element (PE) has available memory spaces in network interface (NI), next, we will send these blocked packets to the local PE, and make the PE re-send these packets to the destination. In order to avoid creating a new circular dependency between the PE and other routers, this operation is carried out only when NI has free memory slots. If the memory in NI is full, we need to drop the current packet and require a higher level end-to-end re-transmission protocol to re-send the packets [10]. For hard link faults, as mentioned in [86], in order to improve network connectivity, we need to selectively remove the turn rules of some routers at the cost of increasing the hazard of deadlock. The scheme discussed above can also be applied to the packets that violate these turn rules so as to return to normal status.

4.4 Simulation Results

4.4.1 Simulation environment setup

In order to evaluate the NoC performance, a systemC based NoC simulator was developed. The energy parameters are adopted from Noxim [79]. The mesh topology is used for the simulation due to its popularity and simplicity. 8×8 meshes are used in the simulation for all the synthetic traffics. Several real benchmarks are also used for evaluating the latency effects. They are MMS (Multimedia system mapped on 5×5 mesh) in [94], DVOPD (Dual Video Object Plane Decoder mapped on 3×4 mesh) and MPEG4 (MPEG4 codec) in [68]. The simulation is carried out for 150000 cycles after 10000 warm up cycles.

4.4.2 Comparison of connectivity of the network

We define fault rate f as follows: $f = \ulcorner \frac{Number\ of\ injected\ faults}{Mesh\ size} \urcorner$; and the parameter α indicates the portion of the partial router faults in the injected fault pattern: $\alpha = \frac{Number\ dynamic\ faults}{Number\ injected\ faults}$; the connectivity of a network is defined as the ratio of

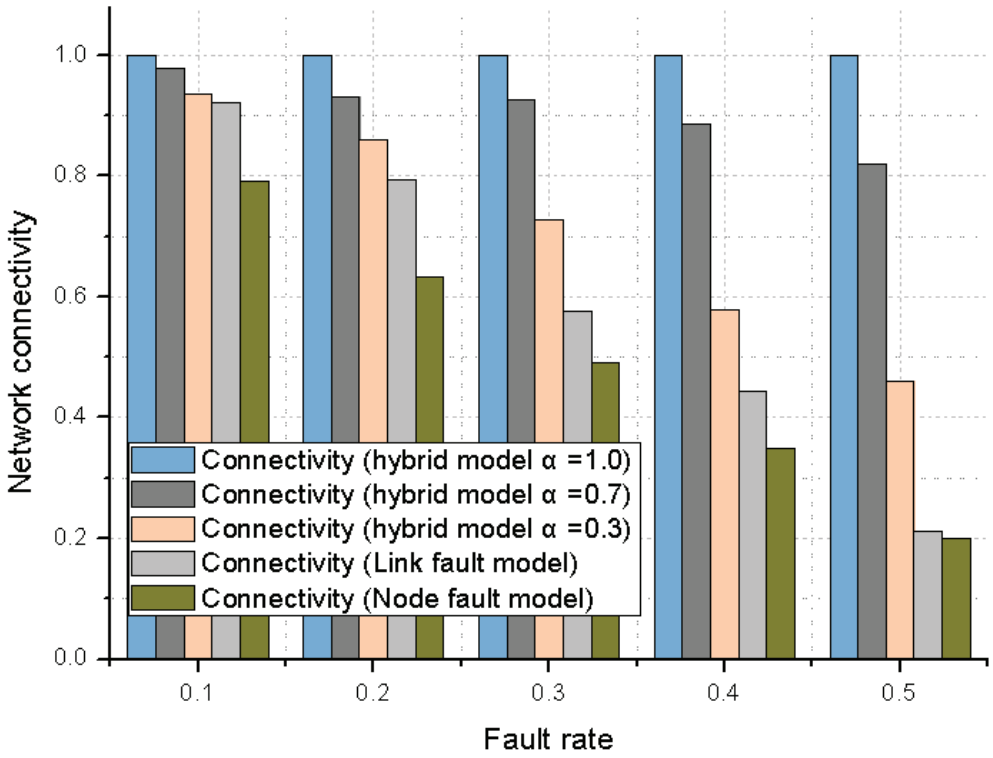

Figure 4.7: Connectivity comparison of the hybrid fault model

number of node pairs that can communicate to total number of node pairs in the network. Figure 4.7 is the analytical results for the connectivity for 8×8 mesh.

For each fault rate, we randomly generate 30 fault patterns. Then the connectivity is analyzed by checking whether each node pair can communicate under the given pattern. From the simulation, we demonstrate by distinguishing the dynamic faults with static faults, network connectivity can be improved more than 20% compared to the link fault model.

4.4.3 Improvement in packet acceptance rate

In table 4.3, we show the simulation results for the packet acceptance rate compared with two state-of-art works in [86, 87]. We randomly generate 20 fault patterns for each fault rate and use random traffic to evaluate this merit. The packets between two nodes which can not communicate anymore will be counted as dropped packets as well. Under high fault rate, the connectivity drops dramatically in [86, 87] and hence reduces the packet acceptance rate. On the contrary, by performing the DBS and DMS operations, the probability of a successful transmission between any two nodes is greatly improved. Hence, 16.6%-21.9% higher

Table 4.3: Packet acceptance rate comparison

Mesh	4×4 mesh					8×8 mesh				
		Previous work		This work			Previous work		This work	
Fault Rate	OE	[87]	[86]	$\alpha = 0.5$	$\alpha = 1.0$	OE	[87]	[86]	$\alpha = 0.5$	$\alpha = 1.0$
0%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
4%	100%	100%	100%	100%	100%	62.3%	100%	100%	100%	100%
12%	55.4%	95.8%	100%	100%	100%	98.2%	100%	100%	100%	100%
20%	33.9%	83.1%	98.3%	100%	100%	25.3%	80.3%	95.9%	100%	100%
28%	29.8%	64.3%	91.0%	100%	100%	13.2%	56.3%	89.6%	93.9%	100%
36%	29.9%	48.9%	74.0%	99.8%	100%	14.2%	32.1%	71.8%	90.4%	96.0%
40%	16.4%	26.0%	69.7%	98.9%	100%	9.7%	23.7%	71.0%	87.7%	92.9%

acceptance rate can be achieved.

4.4.4 Effects on latency performance

In Fig. 4.8, we make a comparison of the proposed round-robin and dedicated DBS operation for 8×8 mesh. Both two strategies outperform the resilient routing in [86] and a $1.15\times$ improvement in the saturation point can be observed.

Fig. 4.9 is the simulation results for 8×8 mesh under the hybrid fault situation. As discussed earlier, DMS will not generate deadlocks so the performance of DMS operation is the closest to the baseline routers. While the DBS operation may bring deadlocks and the latency depends on the predefined threshold of the deadlock detection scheme ($T = 20$ and 64) . It is noticeable that the performance in our work is consistently better than [86] when faults can be categorized as dynamic faults.

In Table 4.4, we summarize the latency and acceptance rate performance for the traffic scenarios other than the uniform random case. Two different fault numbers, 4 and 8 are used for each traffic. We can see that for different traffic scenarios, an average 34.7% reduction in latency and 7.7% improvement in reliability can be achieved.

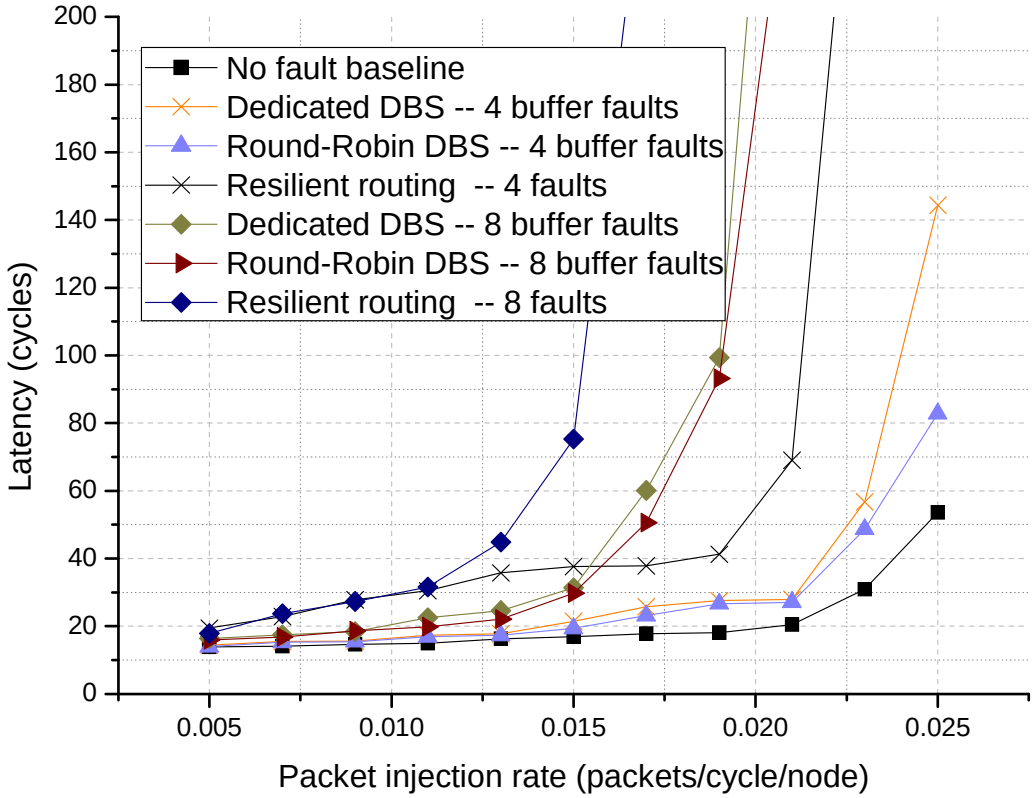

Figure 4.8: Comparison between buffer swapping strategy

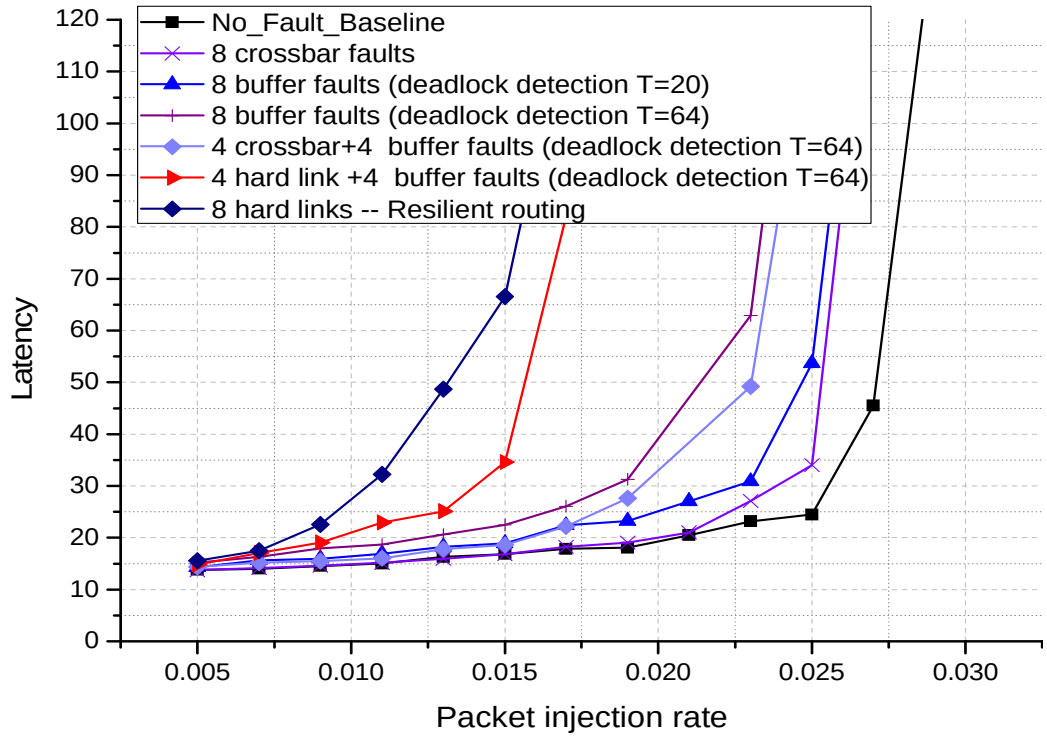

Figure 4.9: Latency comparison of various fault model and schemes

	No fault		Work in [86]		This work		Improvement	
Traffic	Latency (cycle)	Accept Rate	Latency (cycle)	Accept Rate	Latency (cycle)	Accept Rate	Latency (cycle)	Accept Rate
DVOPD-4	18.70	100%	34.24	100%	22.11	100%	35.4%	–
DVOPD-8	18.70	100%	56.22	92.9%	30.91	100%	45.0%	7.1%
MPEG4-4	24.33	100%	70.12	100%	45.51	100%	35.1%	–
MPEG4-8	24.33	100%	82.77	88.6%	47.35	100%	42.8%	11.4%
E3S-4	23.54	100%	27.17	100%	24.87	100%	8.5%	–
E3S-8	23.54	100%	45.33	93.2%	26.20	100%	42.2%	6.8%
MMS-4	39.74	100%	58.18	100%	51.30	100%	11.8%	–
MMS-8	39.74	100%	87.81	94.7%	64.21	100%	26.9%	5.3%
TGFF-4	47.06	100%	66.48	100%	52.15	100%	21.6%	–
TGFF-8	47.06	100%	285.48	100%	64.17	100%	77.5%	–
Average improvement							34.7%	7.7%

Table 4.4: Latency and packet acceptance rate comparison for various traffics

4.5 Conclusion

In this chapter, we propose a fine-grained fault model for network on chips to maximally maintain the network connectivity. Based on the fault model, we distinguish with different fault components and use specific strategies to handle these faults. DBS and DMS techniques are proposed to make the partial router still functional and hence improve the system robustness. Simulation results demonstrate higher packet acceptance rate and better latency performance can be achieved.

CHAPTER 5

FSNOC: A FLIT-LEVEL SPEEDUP SCHEME FOR NOCS USING SELF-RECONFIGURABLE BI-DIRECTIONAL CHANNELS

In this chapter, we tackle the problem of optimizing the bandwidth utilization of the Network-on-Chips (NoC). More specifically, we propose a flit-level speedup scheme to enhance the NoC performance utilizing self-reconfigurable bidirectional channels. For the NoC intra-router bandwidth, in addition to the traditional efforts on allowing flits from different packets using the idling internal bandwidth of the crossbar, our proposed flit-level speedup scheme also allows flits within the same packet to be transmitted simultaneously. For inter-router transmission, a novel distributed channel configuration protocol is developed to dynamically control the link directions. In this way, the bisection bandwidth between the routers can adapt to the changing network status. We describe the implementation of the proposed flit-level speedup NoC on a two-dimensional mesh. An input buffer architecture which supports reading and writing two flits from the same virtual channel at one time is proposed. The switch allocator is also designed to support flit-level parallel arbitration. Extensive simulations on both the synthetic traffic and real application benchmarks, show performance improvement in throughput and latency over the existing architectures using bi-directional channels.

5.1 Introduction

NoCs utilize the interconnected routers instead of buses or point-to-point wires to send and receive packets between processor elements (PE). Due to the high communication bandwidth provided, it overcomes the scalability limitations of the buses and bring significant performance improvement in terms of communication latency, power consumption and reliability *etc* [95]. A typical NoC consists of the processing elements (PEs), network interfaces (NIs) and the routers. The

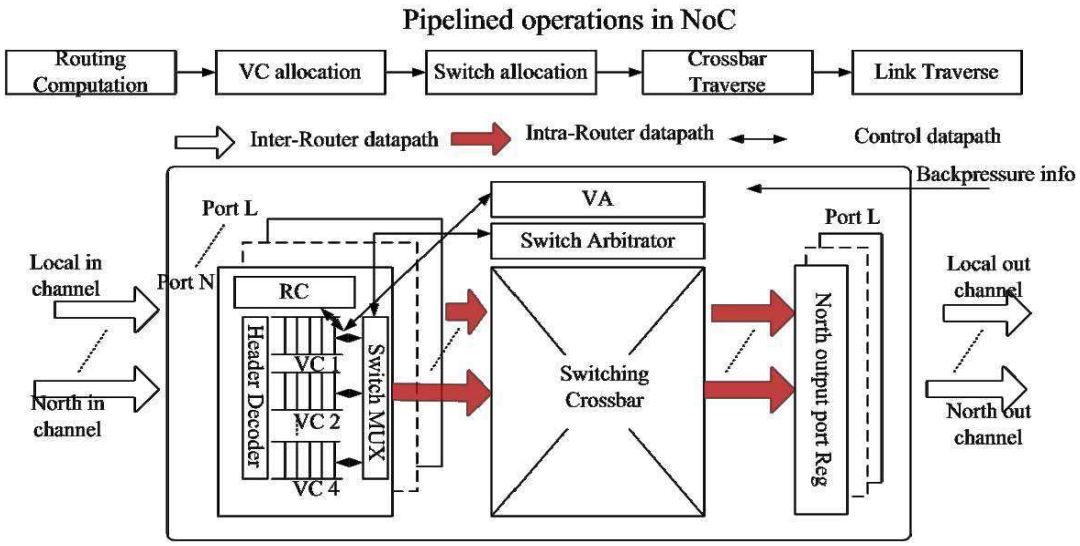

Figure 5.1: Typical four-stage pipelined router micro-architecture using VC-based flow-control

latter two make up of the communication infrastructures in order to support high-bandwidth communication. The NIs are used to packetize the message and transform each packet from the clients (PE cores) into fixed-length flow-control units (*i.e.,* flits). The flits are then sent to the attached router, whose function is to route packets from the input ports to an appropriate output port. In a typical four-stage pipelined, virtual-channel (VC) based router [6] architecture (shown in Fig. 5.1), the router needs to compute the next hop direction (RC stage) and allocates a virtual channel (VA stage) for every incoming header flit. Then the body and tail flits follow the path settled and only need to contend for the switch entries (SA stage). The winning flits will traverse the router to the output port (ST stage) and move to the next hop routers (LT stage). In general, the bandwidth provided by NoC and its utilization have a significant impact on the overall performance [20]. As shown in Fig. 5.1, the NoC bandwidth can be categorized into two types: the inter-router bandwidth (channel bandwidth) and the intra-router bandwidth (switch bandwidth). The inter-router bandwidth is a set of physical interconnection wires between the routers (or router and NI) to handle either the incoming or outgoing traffic from the neighboring routers (or PEs). The intra-router bandwidth is dictated by the datapath within the router (*e.g.,* the input buffers, the crossbar fabric and the output registers).

In order to improve the usage of internal bandwidth, in [20], NoC with input speedup is proposed by providing some excess bandwidth in the cross-point to

relieve the contentions in the routers. In [96], a fine-grained buffer utilization scheme ViChaR is proposed which target dynamic adaption of the router buffers by providing unified buffer resources across all input ports. Due to the efficient usage of the buffers, a 25% performance increase can be achieved over generic routers.

In addition to the efforts on optimizing the internal bandwidth utilization, several approaches have been proposed recently to dynamically adjust the inter-router channel resources to match the run time workload requirements. In [97], based on the observation that NoC link resources are generally over-provisioned, a dynamic channel width configuration algorithm is proposed to adjust the link bandwidth for the current application to optimize the energy-performance trade-offs. In [98], an adaptive physical channel regulator (APCR) is proposed to provide finer granularity of using link channels. The flit size in APCR router is less than the physical channel width (phit). Therefore, the APCR scheme allows flits from different packets to share the same physical bandwidth at run time. However, this concurrent transmission of multiple flits from different packets are implemented at the cost of increased physical link bandwidth. In [99], a heterogeneous NoC architecture is proposed where wider links are allocated to routers that require more resources. At run time, two flits are combined when possible and then sent through the wider link in order to reduce the delay.

Besides the above approaches which are based on the uni-directional links, new progresses have been made by utilizing the self-reconfigurable bi-directional channels. This is motivated by the observation that, between two routers, quite often one uni-directional link may be overflowed with heavy traffic while the link for the opposite direction remains idle [100–102]. The inefficient link bandwidth usage in uni-directional NoC may limit the system latency and throughput. In [100], the link directions are configured in the architecture level with route re-allocation to meet the bandwidth and QoS requirement. In [101, 102], NoC architectures with run time reconfigurable bi-directional channels were proposed. The direction of each link is decided at run time depending on the traffics of the two routers using either an external bandwidth arbiter [102] or a channel direction control (CDC) protocol [101] . For both approaches, the link direction is reversed if there are more than two packets requesting to traverse along one direction and no packets requesting in the opposite direction. Besides the performance improvement in la-

tency and throughput, NoCs with bi-directional channels have also demonstrated the advantages in mitigating the static and dynamic channel failures [103], providing QoS guarantees [104] as well as better energy-performance trade-offs [105].

In this chapter, based on the self-reconfigurable bi-directional channels, we propose a flit-level speedup scheme to further enhance the NoC performance. More specifically, by allowing flits within the same packet to be transmitted simultaneously using the inter- and intra- router bandwidth, the channel utilization is improved and therefore better latency/throughput performance can be achieved. Moreover, a new channel direction control method is proposed which masks inter-router signaling delays within the router pipeline depth and is efficient for configuring the directions of long wire links using repeaters or pipelined registers. The input buffers which support reading and writing two flits from the same virtual channel is proposed. The switch allocator is also designed to support flit-level parallel arbitration. Both the synthetic and real-world traffic patterns have been used to evaluate the performance of proposed flit-level speedup scheme. The hardware overhead is also compared to the conventional NoC and BiNoC designs. The results demonstrate better latency and throughput performance can be achieved with moderate power and area overhead than the traditional NoC and BiNoC schemes.

Compared to the previous works, this chapter makes the following key contributions:

1. A fine-grained, flit-level transmission scheme is proposed which provides more flexibility in using NoC bi-directional channels. We demonstrate the performance enhancement using both synthetic traffic and application benchmarks.

2. A new channel direction control protocol for bi-directional channels is proposed. It is shown this protocol can be used for long wire links with repeaters or pipeline registers.

3. The FSNoC router architecture that supports flit-level parallel transmission is proposed. The performance overhead in terms of area and power are analyzed and compared with state-of-the-art BiNoC implementations.

The remainder of this chapter is organized as follows. Section 5.2 gives the motiva-

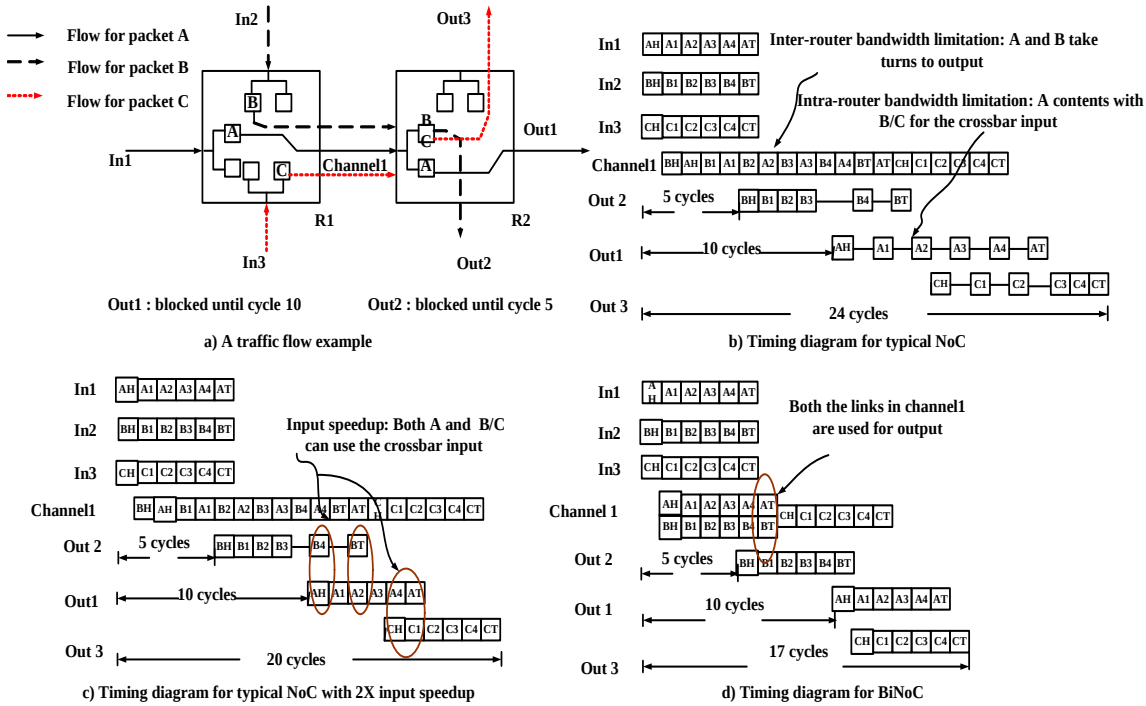
Flow for packet A
Flow for packet B
Flow for packet C
In2
Out3
In1
Channel1
Out1
R1
R2
In3
Out2
Out1 : blocked until cycle 10
Out2 : blocked until cycle 5
a) A traffic flow example
Inter-router bandwidth limitation: A and B take turns to output
Intra-router bandwidth limitation: A contents with B/C for the crossbar input
5 cycles
10 cycles
24 cycles
b) Timing diagram for typical NoC
Input speedup: Both A and B/C can use the crossbar input
20 cycles
c) Timing diagram for typical NoC with 2X input speedup
Both the links in channel1 are used for output
17 cycles
d) Timing diagram for BiNoC

Figure 5.2: The comparisons of different architectures a) example traffic flows b) timing diagram under typical uni-directional NoC c) timing diagram for typical NoC with 2X internal bandwidth d) timing diagram for BiNoCs

tion of adding more flexibility in BiNoC and compares the single router throughput under different transmission schemes. Then, in Section 5.3, we present the proposed flit-level speedup scheme. We elaborate on the channel directional control protocols to support parallel flit transmission in FSNoC, followed by a in-depth discussion of the hardware implementations. Section 5.4 extends the FSNoC design for the existence of long wire links made with repeaters or pipeline registers. Our simulation results are presented in Section 5.5. Finally, section 5.6 concludes the whole work.

5.2 Motivations of FSNoC

5.2.1 A motivational example

To illustrate the difference in bandwidth utilization under different NoC architectures, we consider a case of three packets A, B, C traversing from a router R1 to a router R2 as shown in Fig. 5.2-a. We denote the physical bandwidth of *Channel1* pointing from router the R1 to R2 as B (flits) and the number of input ports in the routers as N. The packets in flow A, B and C have fixed length (in this example it is equal to 6 flits) including the header and the tail

flit. A fair round-robin policy is adopted in the virtual channel and switch arbitration. At the beginning, the packets A, B and C need to contend for the VCs in R2. As there are only two VCs available at each port, we assume packets A and B successfully acquire the virtual channels while packet C has to wait until A or B finishes transmission and releases the occupied VC. We assume the packet stalling [20] occur in R2 during the transmission of packet A and B. The blockage times for packets A and B in ports Out1 and Out2 are 10 and 5 cycles, respectively. Moreover, for the sake of simplicity, in this example, we assume the processing time of a packet transfer is only one cycle in the routers R1 and R2. For the state-of-the-art NoC architectures, there are three ways to utilize the NoC inter- and intra-router bandwidth:

1) *Typical uni-directional NoC switching*: In a typical uni-directional NoC router [6], the internal crossbar size is $N \times N$ while the inter-router bandwidth $B = 1$. Therefore, two kinds of bandwidth limitations exist in the transmission as shown in Fig. 5.2-b. The first limitation is the inter-router bandwidth limitation, where as there is only one unidirectional link between R1 and R2 ($B = 1$), packets A and B have to alternately use this single capacity link (*i.e., Channel1*). The second is the intra-router bandwidth limitation where in R2, the two virtual channels in the west input need to compete one of the N entries in the crossbar during the SA phase. Therefore packet A and B/C can not be transmitted at the same time.

2) *Uni-directional NoC with input speedup switching*: The input speedup scheme provides $2X$ internal bandwidth in the input crossbar side and increases the crossbar size to $2N \times N$. The intra-router bandwidth limitation can be overcome as shown in Fig. 5.2-c, where both the VCs in R2 can use the crossbar entries simultaneously. Thus the latency of packet C is reduced to 20 cycles. However, the inter-router bandwidth limitation still exists in the transmission.

3)*Bi-directional NoC switching*: In BiNoCs [101], since there are no traffic flows from R2 to R1, the link originally pointing from R2 to R1 can be reverted which doubles the inter-router bandwidth from R1 to R2 (*i.e.,* $B = 2$). Consequently, as shown in Fig. 5.2-d, both packets A and B can be transmitted at the same time and the delay of packet C can be further reduced to 17 cycles.

In BiNoC switching, after packets A and B using *Channel1* simultaneously,

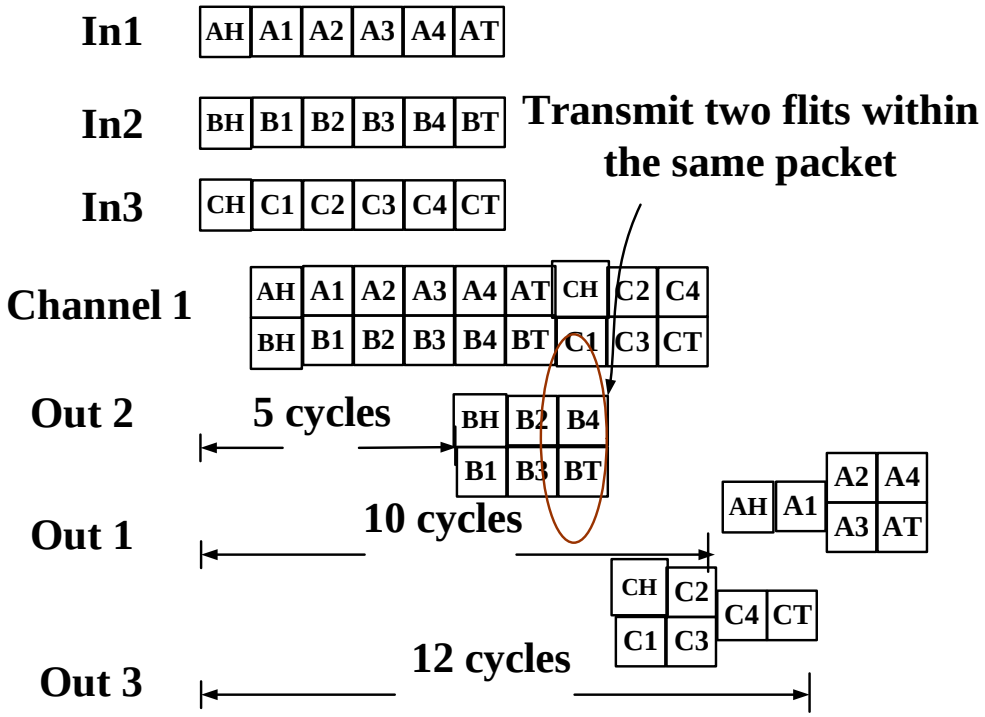

Figure 5.3: The proposed flit-level speedup switching timing diagram for the case in Fig. 5.2-a

the transmission of flits in packet C only needs one flit-width channel (shown in Fig. 5.2-d). In other words, half of the effective bandwidth between R1 and R2 remains unused. Further performance gain can be achieved if we increase the inter-router bandwidth utilization under this case. Towards this end, we propose a flit-level speedup scheme which allows flits within the same packet to use the doubled bandwidth in addition to those from different packets. Fig. 5.3 depicts the timing diagram that allows flits within the same packet (A,B,C) to be transmitted simultaneously. As can be seen from the figure, the latency of packet A/B is reduced due to two flits are sent together towards output port *Out2/Out3*. On the other hand, the latency of packet C is improved because its waiting time is greatly reduced as the packet B in front of it leaves and releases the VC in R2 earlier than in the BiNoC case. In this example, these two effects further improve the latency of the three packets by 12.5%, 27.2% and 29.4%, respectively.

The profiling results in Fig. 5.4 confirms the high opportunity of using the proposed flit-level speedup scheme. More specifically, this is based on the following observations:

1) *The uneven traffic between neighboring routers*: For a channel between two routers, while there is packet sending from one end to the other, there is a

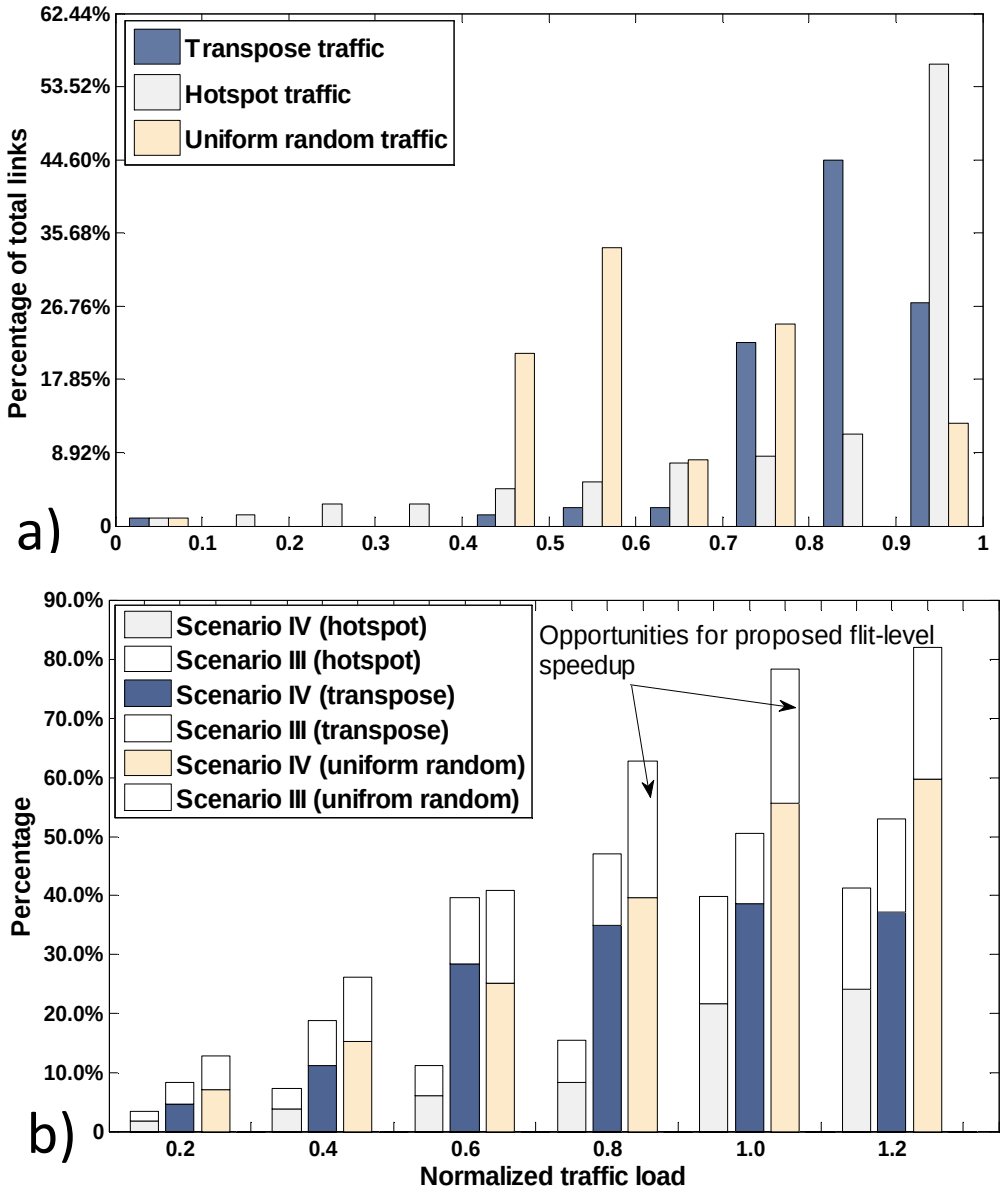

Figure 5.4: Profiling results on an 8×8 mesh under different traffic patterns: a) Histogram of percentage of times that no flows in the opposite link direction at run time b) Breakdown of scenarios III and IV under the condition that there is no contention in the opposite link

significant portion of time that no packet is coming from the opposite direction. Fig. 5.4-a shows the profiling results on an 8×8 mesh NoC architecture under 80% of the saturation injection rate of three synthetic traffic patterns (*i.e.,* uniform random, hotspot and transpose). As can be seen from the figure, in uneven traffic patterns (*e.g.* transpose), for more than 80% of the links, the probability that there is no packets in the opposite direction is larger than 0.7 at run time. Therefore, for these links, more than 70% of time, the effective bandwidth during transmission can be doubled via alternating the opposite link directions. Even for uniform random traffic, most of the links (67.50% in Fig. 5.4-a) have more than 0.5 possibility to double the bandwidth.

2) *The breakdown of the scenarios when there are no flows in the opposite direction:* Under this condition, we can further distinguish into four scenarios for the router under consideration. They are: I) the router has no packet to send through the link, II) the router has only one packet with a single flit to be sent through the link, III) the router has one packet with multiple flits to be sent through the link, IV) the router has multiple packets requesting the output link. For scenarios I and II, we do not need to alternate the opposite link as the original link bandwidth is sufficient for the transmission. However, for the last two scenarios, where scenario III and IV corresponds to flit-level speedup and bi-directional switching, respectively, we can increase bandwidth and hence the performance if we can reverse the direction of the opposite link and send two flits through the two links at the same time. Figure 5.4-b shows the portion of time that scenario III and IV occur among all of the above four scenarios under various injection rates and traffic patterns for an 8×8 mesh NoC architecture. Under low injection rate, since there is little congestion in the network, scenarios I and II dominate the occurrence. On the other hand, when the injection rate increases, the probability that we can use the double capacity of the bi-directional link increases. As shown in the figure, scenario III occurs frequently in run time. However the state-of-art bidirectional switching schemes will only work for scenario IV [101] [102].

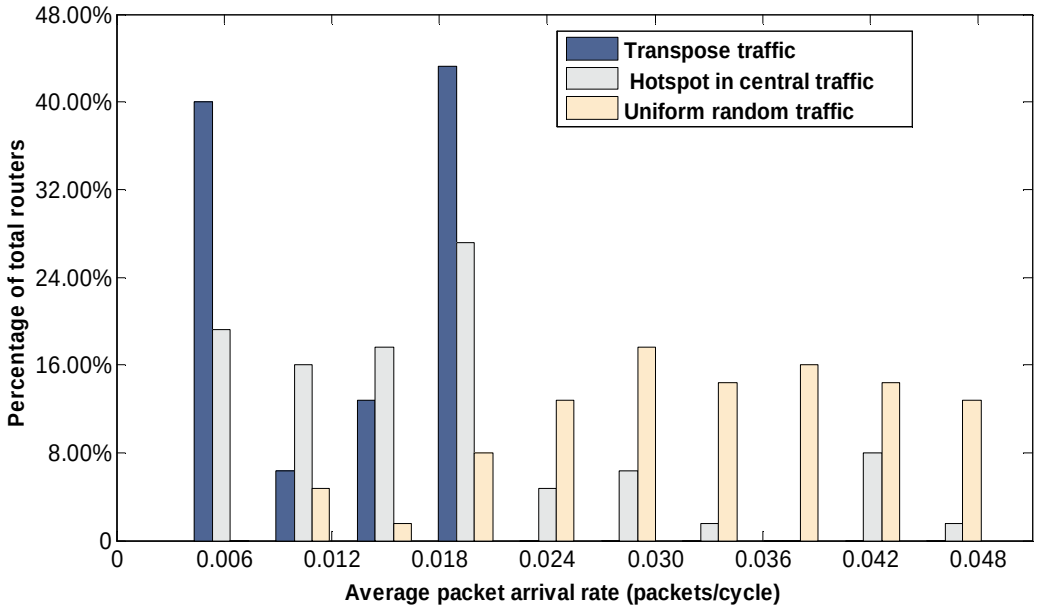

Figure 5.5: The histogram of β on an 8×8 mesh under 80% saturation injection rate

5.2.2 Single router throughput comparisons

FSNoC is further motivated by the throughput comparisons of a single router. We design an event-driven simulator for a single router to model four different switching schemes: 1) Typical uni-directional NoC, 2) Typical NoC with 2X input speedup, 3) Bi-directional NoC and 4) Proposed FSNoC. The simulator is based on two parameters α and β. The α parameter represents the probability that there is no flows in the opposite link (as shown in Fig. 5.4-a) and therefore both two link channels of that direction are enabled for sending data. The β parameter describes the average packet arrival rate at each crossbar input[a]. Both α and β can be obtained by profiling the target applications on specific NoC architectures and are treated as input to the simulator. For example, let β_r denote the mean packet arrival rate of router r over all crossbar inputs in r. For three traffic patterns (random, transpose and hotspot) and under 80% saturation injection rate, the histogram of β_r is shown in Fig. 5.5. As can be seen from the figure, the overall average packet rate β for random, hotspot and transpose traffic patterns are 0.033, 0.013 and 0.018 (packets/cycle), respectively. Similarly, the α values under the same configurations for these three traffics are 0.67, 0.84 and 0.85. These α and β values can be passed to the simulator to evaluate the single router

[a]In a VC-based router architecture, the internal physical bandwidth of a router is mainly dictated by its crossbar input bandwidth. Therefore we abstract the packet arrival process β with respect to each crossbar input entry instead of each VC in the input port.

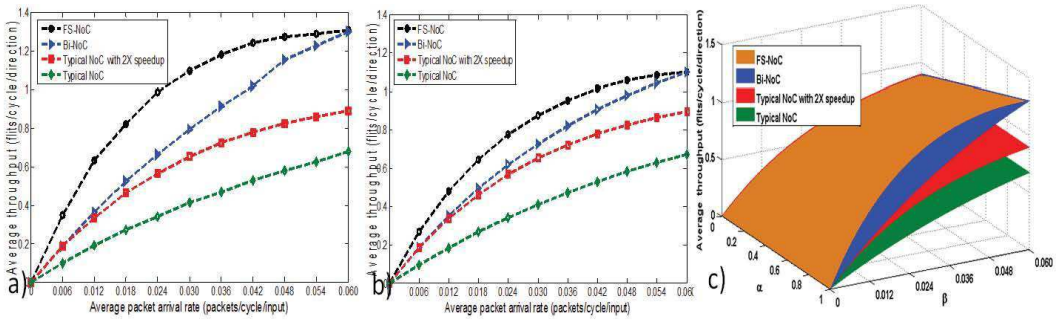

Figure 5.6: Single router throughput comparisons of different router architectures: a) Throughput versus packet arrival rate $\bar{\beta}$ under $\bar{\alpha} = 1.0$. b) Throughput versus packet arrival rate $\bar{\beta}$ given $\bar{\alpha} = 0.5$ c) The throughput comparisons of four NoC architectures under a variety of $(\bar{\alpha}, \bar{\beta})$ combinations.

throughput under a specific traffic pattern. Although being simplified, the single router throughput analysis can provide some useful insights on the performance impacts of FSNoC and other architectures.

In Fig. 5.6-a and -b, The simulated throughputs versus different packet arrival rates $\bar{\beta}$ are shown with $\bar{\alpha} = 1.0$ and 0.5, respectively. It is observed that when the packet arrival rate approximates a very large value (*i.e.*, $\bar{\beta} \rightarrow 0.06$), the BiNoC and FSNoC have the same saturation throughput. This is because under this situation, all the crossbar entries tend to be occupied by different packets and scenario III discussed in sub-section 5.2.1 dominates the time that can use $2X$ bandwidth. However, as profiled in Fig. 5.5, many routers tend to have a smaller $\bar{\beta}$ (*e.g.*, $0.018 - 0.033$) at run time, which provides the opportunities for FSNoC to enhance the system performance. The throughput under a variety of $\bar{\alpha}$ and $\bar{\beta}$ combinations can be obtained by fitting the simulation curves in Fig. 5.6-a and -b as shown in Fig. 5.6-c. As can be observed from the figure, the FSNoC achieves about $10\% - 25\%$ improvement over the BiNoCs for a wide range of α and β values(*e.g.*,$\alpha = [0.6, 1.0]$ and $\beta = [0.012, 0.045]$).

Based on the above observations, we propose FSNoC which improves the NoC performance by adding more flexibility in using the bandwidth. For the inter-router transmission, bidirectional links are employed to provide a double bandwidth at run time. For the intra-router transmission, a new input buffer organization and a switch allocator are designed to allow flits within the same packet to participate in the routing pipelines.

5.3 Implementation of FSNoC

In this section, we present the details of our proposed flit-level speedup scheme. The channel direction control protocol which adapts the link directions based on the run time traffic conditions will be described first. Then we present the router datapath design to support flit-level parallel transmission.

5.3.1 Inter-router channel direction control

In FSNoC with bi-directional channels, both links connecting to one direction of the router can be reconfigured as sending or receiving. However, in order to keep the original order of the flits within the same packet, we define one link as the master and the other as the slave. As shown in Fig. 5.7-a, the master link of a router is also the slave link of its neighbor and vice versa. If there is traffic in both directions between two neighboring routers, the master and slave links acts as the sending and receiving links, respectively. In order to keep the order of flits within the same packet during the flit-level parallel transmission, when writing two flits from the same VC into the output channel, the output controller will always put the first flit on the master link and the second flit on its slave link. On the receiving side, we always assemble the flit appeared on the slave link ahead of the flit on the master link into the VC buffer.

Based on the master and slave link definition, we develop a new channel direction control (CDC) protocol to configure the links at run time so as to adapt for various traffic conditions. Previously, there are two ways to do the direction control, *i.e.,* the centralized and distributed control. For example, [102] employs a separate bandwidth allocator between two adjacent routers so as to determine the link directions simultaneously for both routers, while [101] uses a distributed control mechanism without the need of an extra controller. In this work, the proposed CDC protocol is also based on a distributed scheme since it is more modular and easier to implement. Therefore, we mainly compare the proposed CDC protocol with that used in [101].

For the CDC proposed in [101], there are two Finite State Machines (FSMs) in each router, one for the master link with high priority (*i.e.,* HP-FSM) and the other for the slave link with low priority (*i.e.,* LP-FSM). The direction of each link depends on the state of the corresponding FSM. Since the incoming request

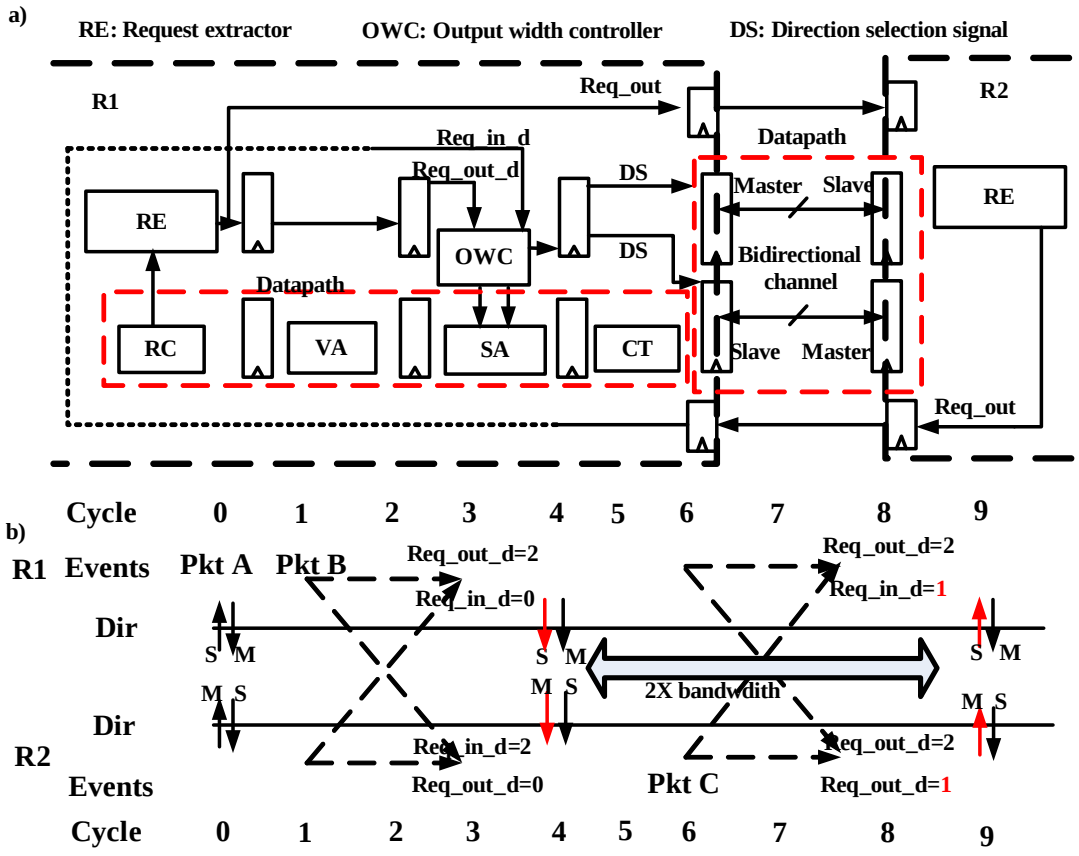

Figure 5.7: a) The channel direction control (CDC) module b) The timing diagram of link direction switching under the arrival of three packets (A,B,C) at router R1 and R2

Algorithm 3 Operation of the RE module

Input: Virtual channel status $VS[N][V]$
Output: Pressure signal $Req\_out[N]$ and local pressure register bank $Req[N]$
if $clock.reset(\,)$ **then**
  for all each input direction $i \in [1, N]$ **do**
    $Req[i] = 0$
    for each VC channel $v \in [1, V]$ **do**
      $VS[i][v].status = idle$;
    end for
  end for
end if{initialize the pressure counter}
always $clock.posedge(\,)$
for all input direction $i \in [1, N]$ **do**
  for all VC channel $v \in [1, V]$ **do**
    if $VS[i][v].status{==}assigned$ **then**
      $k = VS[i][v].Get\_output(\,)$;
      $Req[k].add\_request(\,)$;
      $VS[i][v].status{==}active$;
    else if $VS[i][v].status{==}released$ **then**
      $k = VS[i][v].Get\_output(\,)$;
      $Req[k].decrease\_request(\,)$;
      $VS[i][v].status{==}idle$;
    end if{add or decrease the request}
  end for
end for
for each output direction $o \in [1, N]$ **do**
  if $Req[o].size(\,) \geq 2$ **then**
    $Req\_out[o] = 2$;
  else
    $Req\_out[o] = Req[o]$;
  end if
  $Write\_out(Req\_out[o])$;
end for

signal from the neighboring router takes two cycles to arrive at the current router, a "wait" state is required when the FSM transits from a "receiving" state to a "sending" state which will introduce $1 \sim 4$ dead cycles in the channel direction reversal process [101].

In this work, the proposed channel control scheme is shown in Fig. 5.7-a, which embeds the inter-rouer request signal delays within the four-stages pipeline depth and uses a single module to decide the output width (1 or 2 flits) for each direction instead of two separate FSMs. Two modules, namely Request extractor (RE) and Output width controller (OWC), are introduced to work together with the conventional four pipeline stages in the datapath. The RE module monitors the input channel status and generate the pressure signals to the corresponding OWC modules in the current and the neighboring router. The pressure signal of an output direction, which is named as Req_out in Fig. 5.7-a, represents the number of virtual channels (VCs) that requests to send out towards this direction after being granted a VC in the downstream router. Similar to [101], in order to provide clean signals from router to router in the transmission, we assume both the data and control signals are doubly registered (shown in Fig. 5.7-a). Therefore, it takes two cycles for Req_out to arrive at the neighboring router and become the input request signal (*e.g.,* Req_in_d shown in Fig. 5.7-a) of the OWC module in the neighboring router. The OWC module within each router then works with the switch allocator to process the local pressure signal Req_out_d as well as the received neighbor pressure signal Req_in_d and determines the width of each output direction (2 flits, 1 flit or 0 flit) accordingly. During the crossbar traversal (CT) and link traversal (LT) stage, the direction control signals (DS) of the master and slave link are generated based on the output width calculated in the SA stage and are used to configure the write and read operations.

In the proposed CDC protocol, the pressure signal needs two cycles to arrive at the neighboring router and is required to be used in the SA/OWC stage. In order to guarantee both the OWC modules in router R1 and R2 receive the Req_in_d and Req_out_d signals generated at the same time, the RE module extracts the pressure signals (*i.e.,* Req_out) at the RC stage instead of the VA stage. The effect of generating the signal in the RC stage is that the VC allocation result is updated one clock cycle later when a packet first enters the virtual channel buffer. However, it only affects the performance when the neighboring router requests

to send two flits, while the newly arrived packet in the current router wins the VC allocation after the RC stage. In the next cycle, during the switch allocation (SA) stage, the Req_out_d signal received by the OWC module may still be 0 as it was generated two cycles ago and the latest VA results has not yet been updated. In this case, the current router may treat this as there is no request for sending and the link will then be configured as receiving for both channels while it should be configured as one sending and one receiving. The correctness of the router operation is not compromised, only some of the priority of sending out a flit in current router is lost in this special case. From our simulation results, it is shown that there is no significant difference in the latency and throughput when the pressure signal is generated at the routing computation stage compared with that generated at the VC allocation stage.

Fig 5.7-b shows a typical timing diagram of the link direction switching under the proposed CDC protocol. Assume two packets A and B arrived at $R1$ are granted downstream VCs in cycle 0 and 1, respectively. They are forwarded towards router $R2$. Also, we assume a packet C arrives at router $R2$ and is routed to $R1$ after being allocated a VC in cycle 6. In cycle 1, the RE module in $R1$ extracts a pressure signal $Req\_out = 2$ and sends it to the OWC modules in $R1$ and $R2$. At the same time, since there are no requests from $R2$ to $R1$, the pressure signal generated at $R2$ is still 0. In cycle 3, the OWC module in $R1$ receives the local pressure signal $Req\_out\_d = 2$ from the RE module and the neighbor pressure signal $Req\_in\_d = 0$ from $R2$. The OWC in $R1$ then enables both channels corresponding to the master and slave link during the switch arbitration (SA) for packet A and B. On the other hand, the OWC module in $R2$ disables its master and slave channel during the SA stage to prepare for receiving two flits at one time. After cycle 5, each flit from A and B can traverse the crossbar and the link using the $2X$ bandwidth together. In cycle 8 when the neighboring pressure signal Req_in_out received by $R1$ turns to be 1 due to packet C being granted VC in cycle 6, the slave link channel in $R1$ is disabled in the SA stage and re-configured to the direction of read from cycle 9. In contrast, the master link channel in $R2$ is enabled again during the SA stage for packet C. Therefore, flits from packet C can begin to traverse the crossbar from cycle 9 smoothly.

The operation of the request extractor is described in Algorithm 3. For each VC, its status consists of four states: 1) *idle*: the channel is waiting for the VC

allocation in the corresponding output direction. 2) *assigned*: the packet has just been granted the VC in the downstream router and has not been included in the $Req[N]$ yet. 3) *active*: the VC occupies the downstream VC for transmission. 4) *released*: The tail flit has left the buffer in the last cycle and released the downstream VC. In the beginning, all the channel status are reset to *idle* and the request registers $Req[N]$ are cleared to 0 (line 3-10). At run time, for a specific channel v of direction i, if its status is *assigned*, then we need to get the output direction k of this VC and increased the pressure value $Req[k]$ by one (line 14-17). On the other hand, if the VC status is *released*, we need to decrease the pressure value of direction k by one (line 18-22). Finally, the pressure values in $Req[N]$ are passed to the OWC modules in both current and neighbor router (line 31).

Algorithm 4 describes the operation of the output width controller. It dynamically configures the data width of each output direction (0, 1 or 2 flits) by enabling/disabling the master link ($MLA[N]$) and slave link ($SLA[N]$), respectively. By default, the master link is enabled for write operations while the slave link is disabled for read. At run time, if the local pressure signal ($Req\_out\_d[o]$) indicates that there are 2 requests from different input VCs towards the same direction o and there are no flows in the opposite direction (dictated by $Req\_in\_d[o]$), the controller will enable both the master and the slave link channels corresponding to output direction o to participate in the switch arbitration (line 18-19). Similarly, if the local pressure equals to 1 while the neighbor pressure signal equals to 0, the slave link will be enabled to support flit-level parallel transmission (line 14-16). On the other hand, if $Req\_out\_d[o]$ equals to 0 and the $Req\_in\_d[o]$ is larger than 0, both the master and slave links will be disabled to let the neighboring router output using the $2X$ bandwidth (line 12-13). In addition to generate MLA and SLA signals, the OWC module also determines the input switch arbitration mode $SA\_in[N][V]$. By default, $SA\_in[N][V]$ for each VC is set to be *normal* (line 4-8), which means only one flit from this VC can participate in the SA stage every cycle. However, under the case of flit-level speedup, the VC can send two flits at the same time. Therefore, under this condition, the switch arbitration mode of the requesting VC is set to be $flit\_speedup$ (line 17), which allows two flits within the same packet to participate in the arbitration for both the master and slave links of the corresponding output direction.

Algorithm 4 Operation of the OWC module

Input: Pressure signal $Req\_out\_d[N]$ and $Req\_in\_d[N]$
Output: Master/slave link enable $MLA[N]$ and$SLA[N]$, Switch allocation mode $SA\_in[N][V]$
always $clock.posedge(\ )$
for all input direction $i \in [1, N]$ **do**
  for all VC channel $v \in [1, V]$ **do**
    $SA\_in[i][v] = normal$;
  end for
end for
for all output direction $o \in [1, N]$ **do**
  $MLA[o] = enable$;
  $SLA[o] = disable$;
  if $Req\_out\_d[o] == 0 \&\& Req\_in\_d[o] > 0$ **then**
    $MLA[o] = disable$;
  else if $Req\_out\_d[o] == 1 \&\&\ Req\_in\_d[o] == 0$ **then**
    $SLA[o] = enable$;
    $(i, v) = get\_input\_vc(o)$;
    $SA\_in[i][v] = flit\_speedup$;
  else if $Req\_out\_d[o] == 2\ \&\&\ Req\_in\_d[o] == 0$ **then**
    $SLA[o] = enable$;
  end if
end for

5.3.2 Router datapath design of FSNoC

Besides the RE and OWC modules added in the control path, the router datapath also needs to be modified. In FSNoC, the routing computation (RC) module and virtual channel allocator are identical to those in the typical NoC [106] and BiNoC architecture [101] because the output directions/VCs are computed/reserved in a same way in all architectures. However, the FSNoC has a number of additional requirements in the input buffer and switch allocator design to support efficient flit-level parallel transmission.

Input buffer organization

In the proposed flit-level speedup scheme, it is possible that two incoming flits are for the same VC buffer. Therefore, it is required to support reading/writing two flits from/to the same VC buffer at the same time. One way is to adopt multiple ports memory. However, this causes significant overhead in terms of area and memory access time [20]. Here we propose a novel input buffer organization, which is shown in Figure 5.8, to satisfy this requirement.

Figure 5.8-a shows the input port configuration with V virtual channels. Two

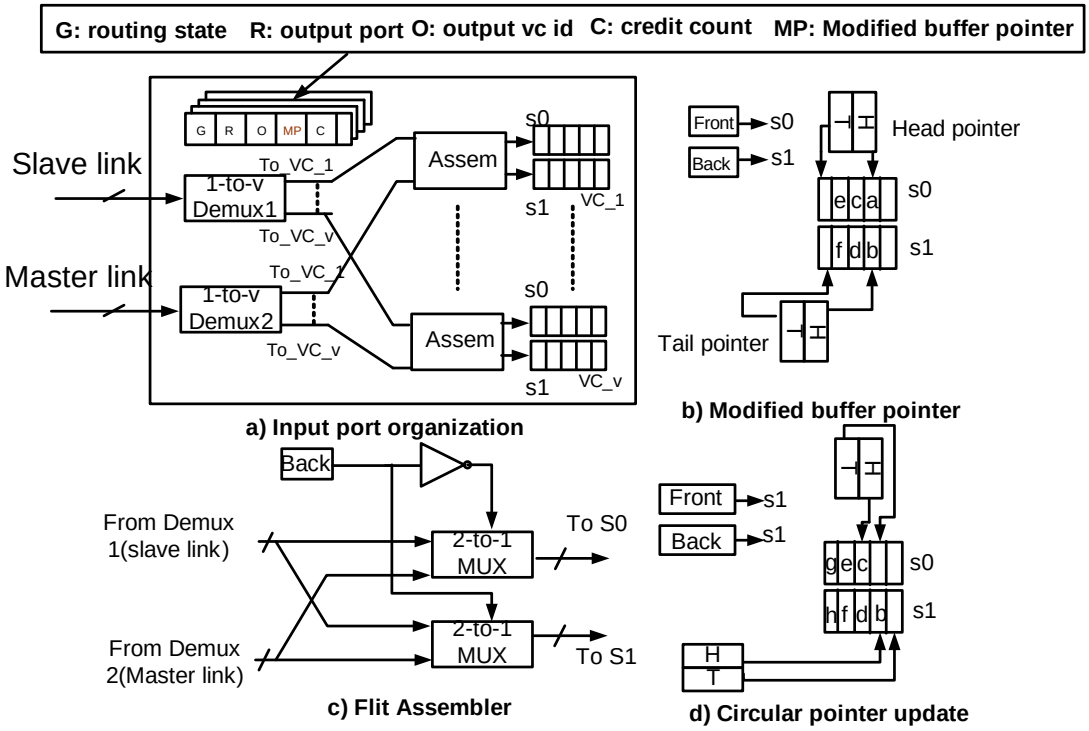

Figure 5.8: Proposed buffer organization

$1 - to - V$ demux and V flit assemblers are needed for each port. The original virtual channel buffer is spitted into two sub-buffers, s0 and s1, respectively. Both sub buffers share the same buffer state information (the GROPC vector shown in Figure 5.8-a are defined as in [20]). The incoming flits are written into s0 and s1 alternatively so that two consecutive flits are stored in different sub buffers. Figure 5.8-b shows the modified buffer pointer in each virtual channel. For s0 and s1, each has a head-tail pointer pair pointing to the start and end addresses of the flits of the packets. In addition, two 1-bit registers, denoted as the Front and Back registers, are added to indicate which sub-buffer holds the first and the last flit (*e.g.,* flits "a" and "f" in Figure 5.8-b) of the VC channel.

The details of the flit assembler is described in Figure 5.8-c. It is responsible for assembling the flits into the sub-buffer of the VC in the right order. When two flits are sent in the reconfigurable links, the first flit is always connected to the master link of the sender (*i.e.,* the slave link of the receiver). To ensure the correct assembling of the flits, the slave link is always connected to the sub-buffer that is not pointed by the Back register. Figure 5.8-b shows an example. The Back register points to s1 since the last flit "f" is stored in s1. The assembler will connect the slave link to s0 and the master link to s1 in the next operation.

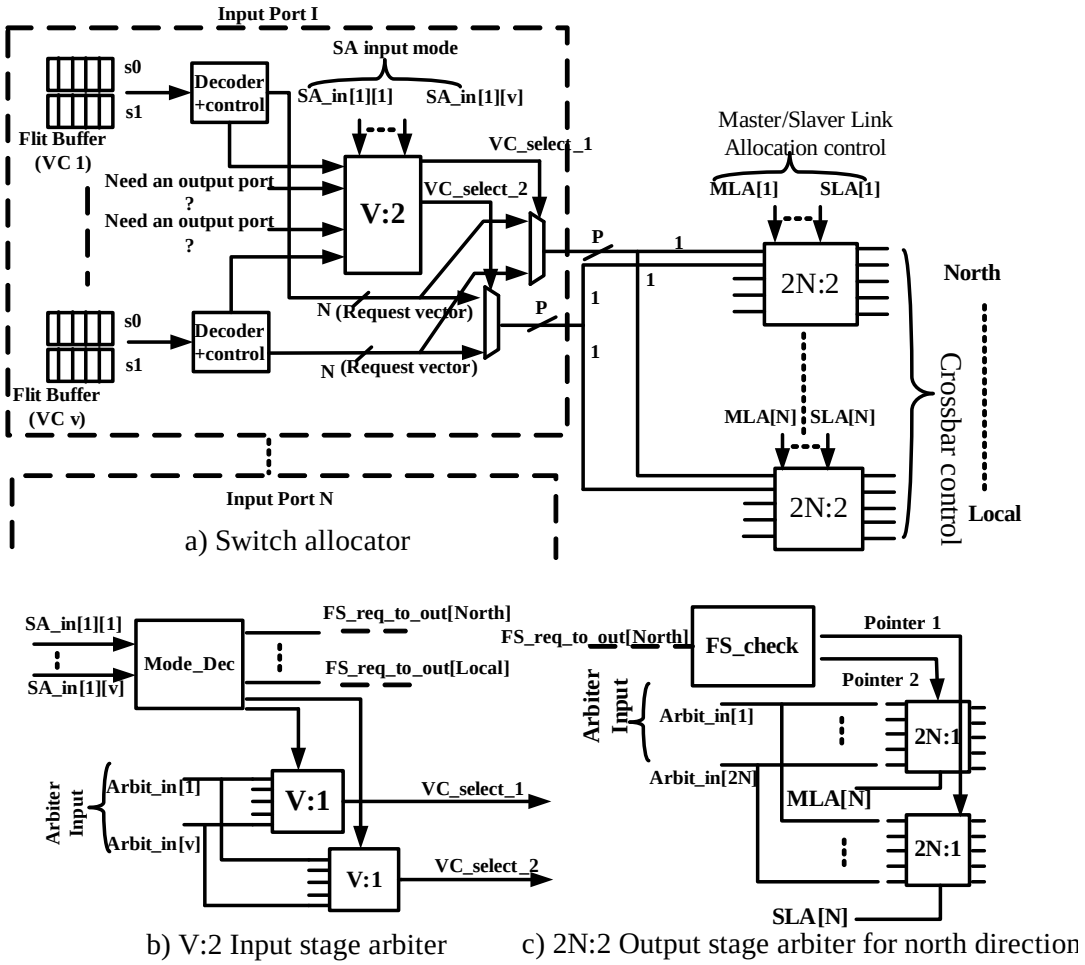

Figure 5.9: Switch allocator logic (modified from [106])

Similar operation is executed when only 1 flit is written into the buffer from the input link. The read operation is similar to the write operation based on the value of the Front register. Figure 5.8-d illustrates the updated pointers after reading a flit "a" from the buffer and writing flits "g" and "h" into it. As shown in the figure, the head and tail pointers are updated in a circular manner while the Front and Back registers are updated based on the number of flits reading from or writing into the virtual channels.

Switch allocator design

The switch allocator assigns the master and slave link channel of each output direction to the input VCs such that the VCs with granted switch access can move flits to the corresponding links in the crossbar and link traversal stage. Different from BiNoC, to support the flit-level speedup in the NoC, the switch allocator should not only support granting two requests from different VCs, but also allow

a single VC to arbitrate for both the master and slave links. Figure 5.9-a shows the building block of our switch allocator. This allocator is accomplished with a two-stage arbitration architecture [20], which is shown in Fig. 5.9-a. At the input stage, a $V : 2$ arbiter in each input direction is used to select 2 requests out of many requests from the V virtual channels. Of note, the 2 winners can come from different VCs or from the same VC. This is determined by the SA arbitration mode vector ($SA\_in[N][V]$) discussed in section 5.3.1. If all the VCs in the current input direction i have the same mode *normal* (*i.e.,* $SA\_in[i][v] = normal$, $\forall v \in [1, V]$), then the two winning requests come from different VCs. On the other hand, if some of the VCs have the SA mode $flit\_speedup$, the two winning requests are combined together and returned to one of these flit-level speedup VCs. At the output side of the switch allocator, N $2N : 2$ arbiters are utilized to decide the actual number of the requests that is finally granted.

Figure 5.9-b shows the detail design of the $V : 2$ arbiter for the input stage allocation. It consists of two $V : 1$ round-robin arbiters and a mode decision module named as $Mode\_Dec$. The $Mode\_Dec$ module of direction i reads in the SA mode vector (*i.e.,* $SA\_in[i][v]$, $v = [1, V]$) and determines the selection mode for the current input port. If all the entries of SA vector equal to *normal*, it means that there is no request that requires two flits to be transmitted from the same VC and hence the two $v : 1$ arbiters will work independently to select two VCs using different priority pointers . Otherwise, the $Mode\_Dec$ will chooses one VC with the $flit\_speedup$ mode and generate the priority pointer exactly pointing to that VC for both $V : 1$ arbiters.

Figure 5.9-c shows the $2N : 2$ arbiter used in the output side. It consists of two $2N : 1$ arbiters. The link allocation control signals ($MLA[N]$ and $SLA[N]$ discussed in Algorithm 4) of each output direction enable/disable the arbiters accordingly. Moreover, a module named as FS_check (shown in Figure 5.9-c) is used to grant two requests to the same input port so to allow parallel transmission of two flits from the same VC.

5.4 FSNoC for long-wire links

For many NoC designs, there exist long wire links due to the irregular topology [107] or the short-cut paths created by long-range links [108]. In order to address

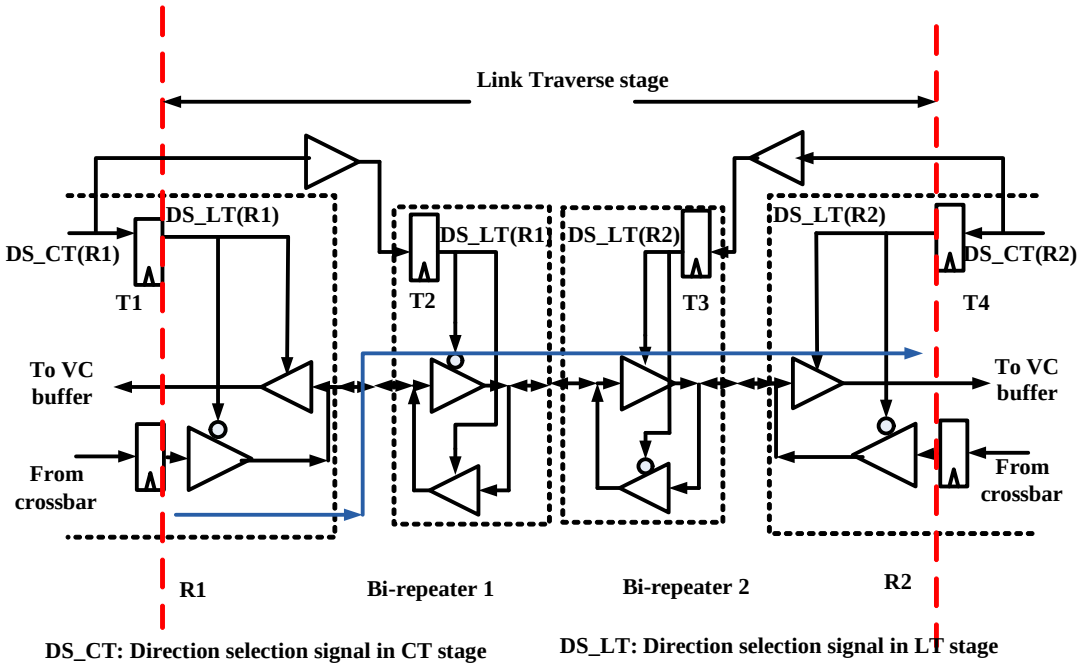

Figure 5.10: FSNoC design for long wire link with repeaters

the issue of high latency and power consumption, repeaters [109] or pipeline registers [107, 108] are usually inserted between the routers to break the long wires. In this section, we extend the FSNoC design to support flit-level parallel transmission over the long wires.

5.4.1 Long-wire links with repeaters

For long-wire link with repeaters, it still takes one cycle to move a flit from one router to the neighbor. In FSNoC, the control channels (*e.g., Req_out* in Fig. 5.7-a) between two routers are uni-directional. Therefore, the repeaters can be inserted directly between the sender and receiver as in the case of typical NoCs. On the other hand, for the data channels, bi-direction repeaters (Bi-repeater) are needed to allow the data transmission over both directions. Moreover, the direction selection signal (*i.e.*, DS in Fig. 5.7-a) should be propagated to configure the intermediate Bi-repeaters properly. Fig. 5.10 shows the detail design of a data channel between two routers $R1$ and $R2$. Within each router, the OWC module determines the output channel width as well as the DS signal in the SA stage based on the CDC protocol. Then, the generated DS signal is pipelined together with the flits in the crossbar traverse (CT) and link traverse (LT) stages before reaching the output port to configure the inout directions. As shown in

Fig. 5.10, the DS signal during the CT and LT stages are represented as DS_CT and DS_LT, respectively. If there are n Bi-repeaters need to be inserted between router $R1$ and $R2$ to address the long wire delays (Seen Fig. 5.10 for a case of two Bi-repeaters), $\lfloor \frac{n+1}{2} \rfloor$of them close to $R1$ will be driven by the DS signal from $R1$ while the remaining are driven by the DS from $R2$. For example, in Fig. 5.10, Bi-repeater 1 and 2 are driven by $R1$ and $R2$, respectively. The one-cycle ahead signal DS_CT from $R1$ and $R2$ are connected to the registers in the Bi-repeater 1 and 2. Every cycle, at the clock rising edge, the DS_CT signals will be read into the registers $T1 - T4$ simultaneously. They become the DS_LT signals for the current cycle. Then, the DS_LT signals configure the Bi-repeater to ensure the write/read directions are the same as the driving router. As the CDC control protocol discussed in Section 5.3 guarantees $DS\_LT(R1)$ and $DS\_LT(R2)$ are conflict-free, the correct data transfer direction can be maintained in the Bi-repeaters. Also, as the number of pipeline stages do not change, the FSNoC with repeaters has the same latency as the original design for data transmission.

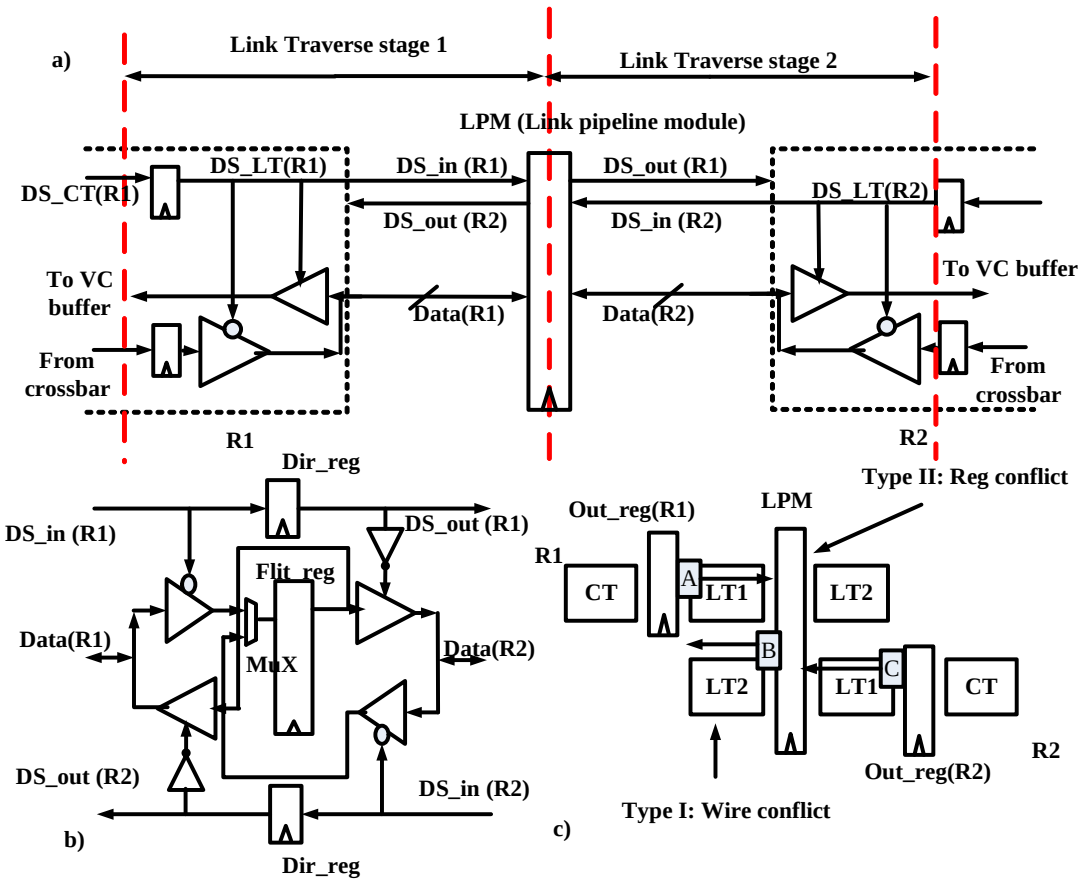

Figure 5.11: a) FSNoC design for long wire links with pipeline registers, b) the detail schematics for data link pipeline registers, c) two types of conflicts during the the link traversal (LT) stage

5.4.2 Long-wire links with pipeline registers

Besides repeaters, pipeline registers can also be used to break long wire links which achieves better timing properties at the cost of longer link traversal latency.

For the control channels, conventional registers can be added directly as the signals are uni-directional. On the other hand, for the data channels, the direction selection signals (DS signals) are pipelined together with the data flits to ensure the proper read and write direction. This is shown in Fig. 5.11-a, where a link pipeline module (LPM) is inserted between routers $R1$ and $R2$ for a data link. Two direction selection signals from $R1$ and $R2$, *i.e.,* $DS\_in(R1)$ and $DS\_in(R2)$, are also connected to the LPM. The detail schematic of the pipeline module is shown in Fig. 5.11-b. In LPM, a flit register ($Flit\_reg$) and two one-bit registers (Dir_reg) are used to buffer the data flits and direction signals, respectively. Every cycle, based on $DS\_in(R2)$ and $DS\_in(R1)$, a MuX selects whether the flit on $Data(R1)$ or $Data(R2)$ should be written into the internal $Flit\_reg$. Similarly, at the output side of $Flit\_reg$, the previous one-cycle pipelined DS signals, *i.e.,* $DS\_out(R1)$ and $DS\_out(R2)$, determine the link segment ($Data(R1)$ or $Data(R2)$) that current flit in $Flit\_reg$ needs to be forwarded to.

Compared to FSNoC in section 5.3, the additional link traversal cycles due to link pipelining complicate the channel direction control in two aspects. First, it brings the *synchronization* issue between the neighboring and local pressure signals. As shown in Fig. 5.7-a, originally, it takes two cycles for both the neighbor pressure (Req_in) and local pressure signal (Req_out) to arrive at the OWC module. As the OWC modules in two adjacent routers make the decision based on same (Req_out_d, Req_in_d) pair every cycle, it is ensured that no write conflicts (*i.e.,* write from both directions) occur by the proposed distributed CDC protocol in Algorithm 4. On the other hand, here, for the case of K-stage pipelined link, the pressure signal Req_out takes additional K cycles (*i.e.,* $K+2$ cycles) to arrive at the OWC in $R1$ while it still needs two cycles to reach OWC in $R2$. The asynchronization of OWC modules in $R1$ and $R2$ may create incorrect output width combinations (*e.g.,* an output width of 2 flits for $R1$ and 1 flit for $R2$) and overwrite useful flits on the link. Therefore, in order to provide correct direction decisions, K additional pipeline registers should be added between the RE and OWC modules inside each router to keep the local pressure signals synchronized

with the neighbor pressures.

The second hurdle due to link pipeline is that *more types of conflicts* need to be resolved in the CDC protocol. More precisely, two types of conflicts on a pipelined link are identified, namely *wire conflict* and *register conflict*. Taking a one-stage pipelined link in Fig. 5.11-a and -b as an example, the *wire conflict* happens if both ends try to write data on the same wire segment. As shown in Fig. 5.11 -a and -c, the output register of router $R1$ will write a flit A on the link segment $Data(R1)$ if the direction selection signal $DS\_in(R1)$ equals to OUT (*i.e.,* 0). On the other hand, the LPM will write a flit B on the same link segment if the direction selection signal $DS\_out(R2)$ equals to OUT. Therefore, a wire conflict happens if both $DS\_out(R2)$ and $DS\_in(R1)$ equal to OUT at the same time. The *register conflict* refers to the operation that try to write flits into the same LPM in one cycle. For example, if $DS\_in(R1)$ equals to OUT, flit A on $Data(R1)$ will be written into LPM at the next cycle. Similarly, if $DS\_in(R2)$ equals to OUT, flit C on $Data(R2)$ will be written into LPM. Thus, the register conflict occurs if $DS\_in(R1)$ and $DS\_in(R2)$ equal to OUT at the same time. As discussed in section 5.3, in FSNoC, the direction selection signal $DS\_in(R1)$ and $DS\_in(R2)$ are generated in the same cycle by the OWC module of $R1$ and $R2$, respectively. Therefore, they cannot be OUT simultaneously under the CDC protocol in Algorithm 4, which means the *register conflict* will not occur. On the other hand, as $DS\_out(R2)$ is generated and pipelined one cycle before $DS\_in(R2)$, they may both be equal to OUT at run time and produce the *wire conflicts* as a result. In order to resolve the wire conflicts in one-stage pipelined link, the OWC module needs to store the direction selection results in the previous cycle and ensure the DS signal generated in current cycle not conflict with that of the neighbor router in both current and previous cycles. More generally, for a K-stage pipelined link with $K \geq 2$, more than one LPMs are inserted between two routers. In order to read and write correctly, we need to guarantee there are no wire conflicts on the $K+1$ link segments as well as no register conflicts in the K LPM modules. Therefore, the OWC module in each router needs to store the DS signals of the previous K cycles. For the master link, its direction will be reversed for the other router if the current router have not written flits into this link during the past K cycles. Similarly, for the slave link, its direction can be reversed for write operations if the current router receives the

Table 5.1: NoC architectures used in experiments

Architecture	Buf. /Dir. (flits)	Buf. /VC (flits)	Ch./Dir.	Crossbar size	Freq. (MHz)	Norm. T
Typical NoC	64	16	1-*in* 1-*out*	5×5	1265	1.0
2X speedup	64	16	1-*in* 1-*out*	10×5	1219	1.04
BiNoC	64	16	2-*inout*	10×10	1020	1.24
FSNoC	64	16	2-*inout*	10×10	1020	1.24

neighboring pressure signal $Req\_in = 0$ continuously during the past K cycles.

5.5 Experimental results

5.5.1 Simulation setup

We evaluate and compare the proposed FSNoC architecture with other designs using a cycle-accurate, SystemC-based simulator extended from Noxim [79]. We assume the mesh-based topology for all the comparisons. Table 5.1 shows four router architectures used for comparison. Specifically, we compared the performance of the proposed scheme with three different architectures, namely the Typical NoC, Typical NoC with 2X input speedup, and BiNoC. For all architectures, each input direction of the router has 4 VCs and the buffer depth of each VC is 16 flits as in [101]. We assume the packet has a constant length of 16 flits. As shown in Table 5.1, for conventional unidirectional NoC, a 5×5 crossbar is used to support the 5-in and 5-out data transmission. The input speedup scheme provides $2X$ excess bandwidth in the input side and therefore, a 10×5 crossbar is needed. For the BiNoC and FSNoC architectures, as we need to provide flexible data transmission over the 10 inout ports, a larger 10×10 crossbar is required. In order to make a fair comparison over the four architectures, each router architecture is assumed to operate under the maxium frequency, which is obtained from the synthesis results by Synopsys Design Compiler under TSMC 65nm technology. The simulated latency is then calculated in the unit normalized to the cycle T of Typical NoC.

For four architectures, the Dimension ordered XY routing is adopted in order to avoid deadlocks (except Fig. 5.14-a, which evaluates the effects of oddeven

routing algorithm). Various traffic patterns were used in the evaluation, including synthetic traffic and real benchmarks. More specifically, two types of real benchmarks are used in the evaluation. The first type is task-graph-based MPSoC applications, including MWD (Multi-Window Display) [68], MMS (Multimedia system) [67], MPEG4 (MPEG4 codec) [68] and DVOPD (Dual Video Object Plane Decoder) [62] as well as three E3S [110] applications named auto_indust, telecom and consumer. For these application-specific benchmarks, we need to map the task graphs onto the mesh architecture first. The second type of applications are extracted from SPEC-web [111] benchmarks which model the memory access patterns for CMP architectures. For SPEC-web applications, five 16-node multithreaded workloads for IBM and Oracle database server (*i.e., DB2V* and *Oracle*), the Apache HTTP server (*i.e., Apache*), scientific workloads of matrix factorization (*i.e.,Sparse*) and the ocean dynamic simulation (*i.e., Ocean*) are extracted on a 4×4 mesh. Moreover, application consolidations [112] are evaluated by choosing four applications and randomly mapping them onto an 8×8 mesh (*i.e., Mixed*). In the following, we first evaluate the performance of the NoC architectures without link pipeline stages in sub-section 5.5.2-5.5.3. Then, we discuss the simulation results for FSNoC designs link pipelines.

5.5.2 Simulation results on synthetic traffics

For the synthetic traffic comparison, six traffic patterns were used, *i.e.,* random, transpose, hotspot, shuffle, bit-reversal and butterfly [20, 79]. For the hotspot traffic, two central nodes in the mesh network have 10% higher probability to receive packets from other nodes. Figures 5.12 summarizes the comparisons on the latency under the XY routing algorithm. For all the cases, the proposed FSNoC out-performs other three schemes. As observed from the figure, the improvement highly depends on the traffic patterns. For example, comparing the results in random traffic with transpose or bitreversal traffic, we can find the uneven traffics achieve more latency improvement as these non-uniform patterns provide more chances for the BiNoC and FSNoC architecture to use the $2X$ bandwidth during the transmission. Specifically, we can compare the maximum packet injection rate sustainable by the network of each architectures [23]. The maximum injection rate (*pir*) is calculated as the rate (packets/cycle/node) at which the corresponding throughput reaches 95% of the saturation. For instance, for the random traffic

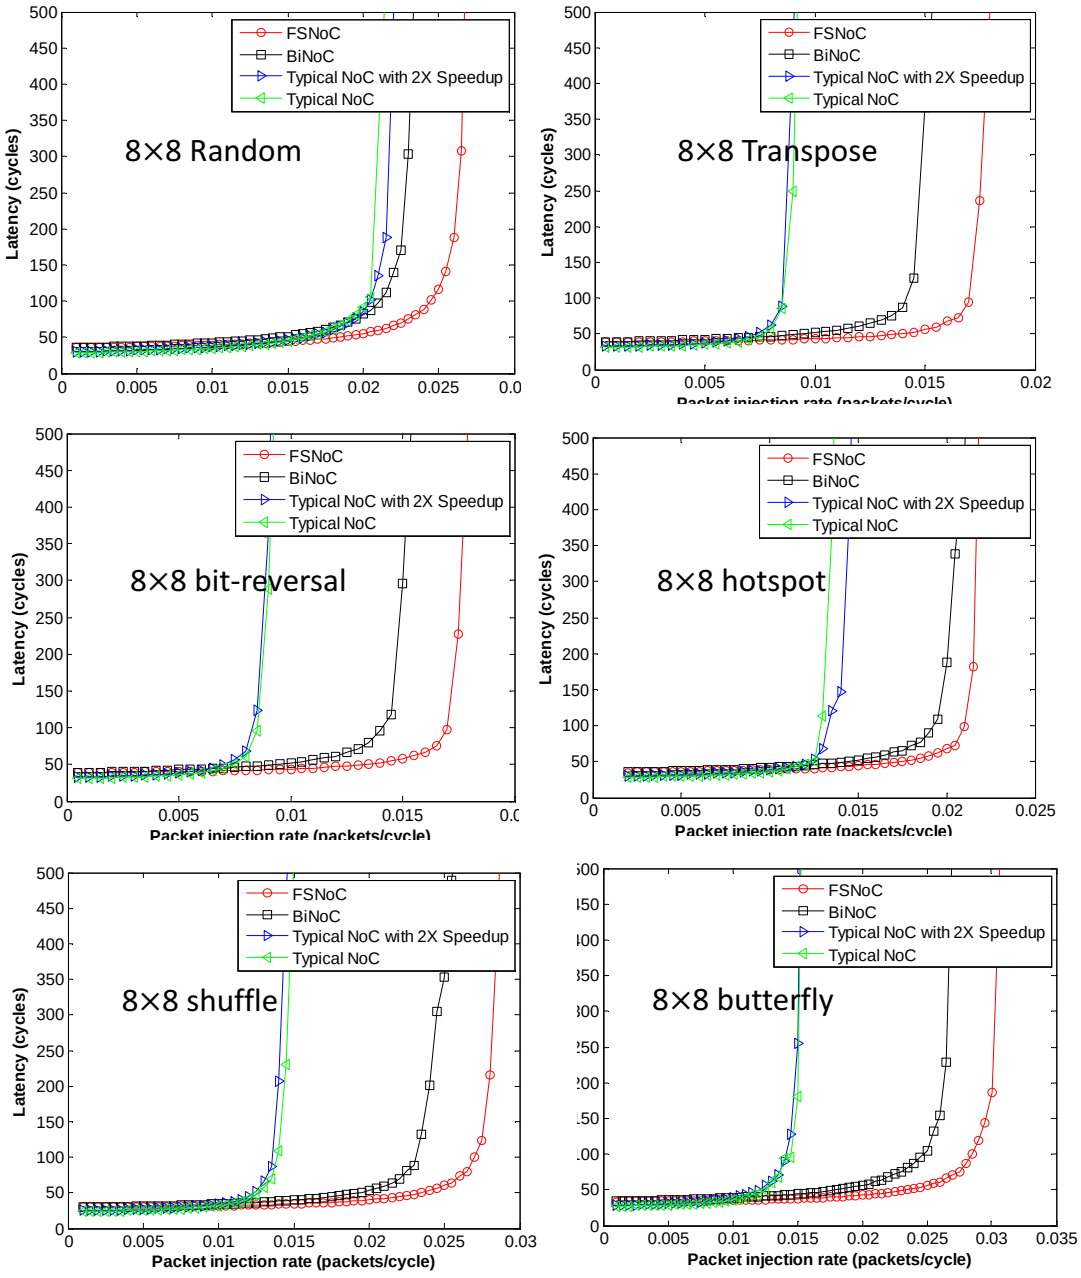

Figure 5.12: Simulation results for six synthetic traffics on an 8×8 mesh under XY routing

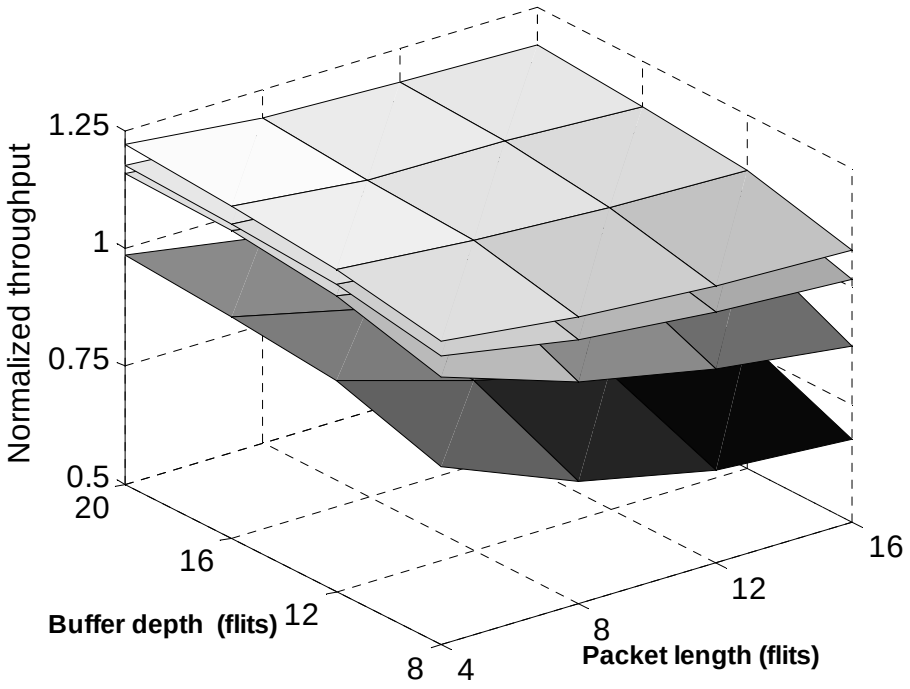

Figure 5.13: The impacts of buffer depths and packet size on overall throughput

pattern, the *pir* of FSNoC, BiNoC, $2X$ speedup NoC and Typical NoC are 0.020, 0.021,0.022 and 0.0255 (*packets/cycle*), respectively. While for transpose traffic, the *pir* of four architectures are 0.008, 0.008, 0.014, 0.016 (*packets/cycle*). In summary, the BiNoC architecture improves *pir* by $10\% - 72\%$ over the typical and $2X$ input speedup schemes. In addition to that, the FSNoC further increases the saturation injection point by $12\% - 22\%$.

We then evaluate the NoC performance regarding the influence of different packet length ($4 - 16$ flits) and buffer depth ($8 - 20$ flits per VC). Figure 5.13 shows the maximum throughput that four architectures can obtain for the random traffic pattern on an 8×8 mesh. All the throughput values are normalized to that of a typical NoC design with constant packet length of 4 flits and buffer depth of 20 flits. As shown in the figure, the throughput improvement is more obvious when a longer packet length is adopted as the single packet transfer time can be more significantly reduced by allowing sending two flits at one time. Overall, the proposed FSNoC design achieves the highest throughput over a wide variety of packet length and buffer depth combinations.

In Fig. 5.14-a, we show the latency of four architectures running oddeven routing algorithm under random and transpose traffic patterns. For the uneven

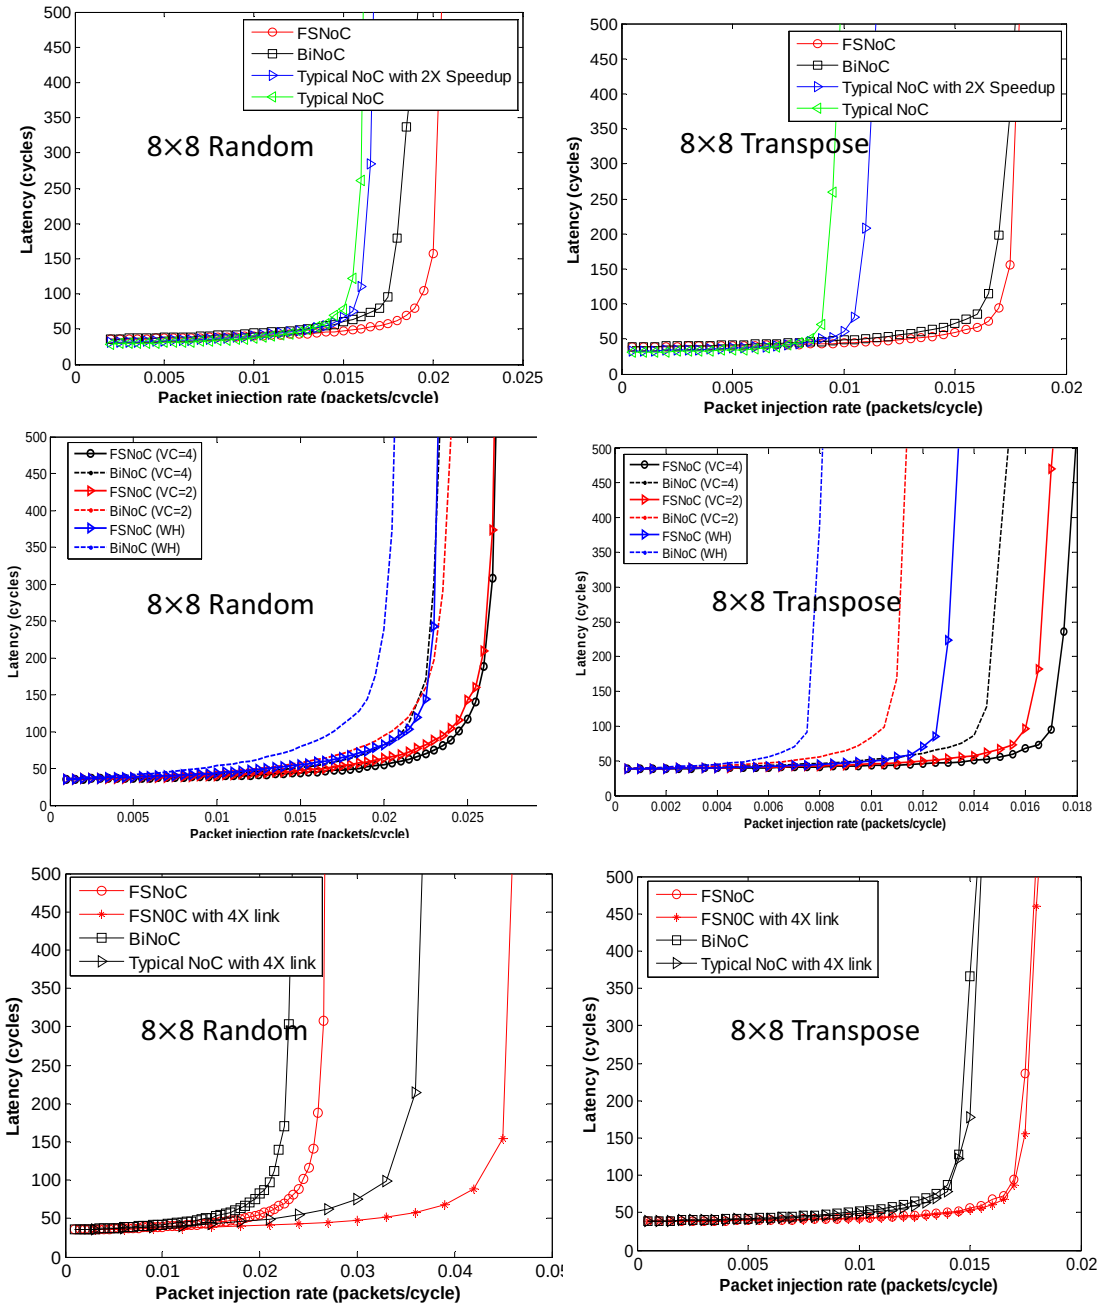

Figure 5.14: Latency comparisons for a) four architectures under oddeven (OE) routing algorithm b) FSNoC and BiNoC with WH router architecture and VC router architecture with different VC numbers c) NoC architectures with 4L links

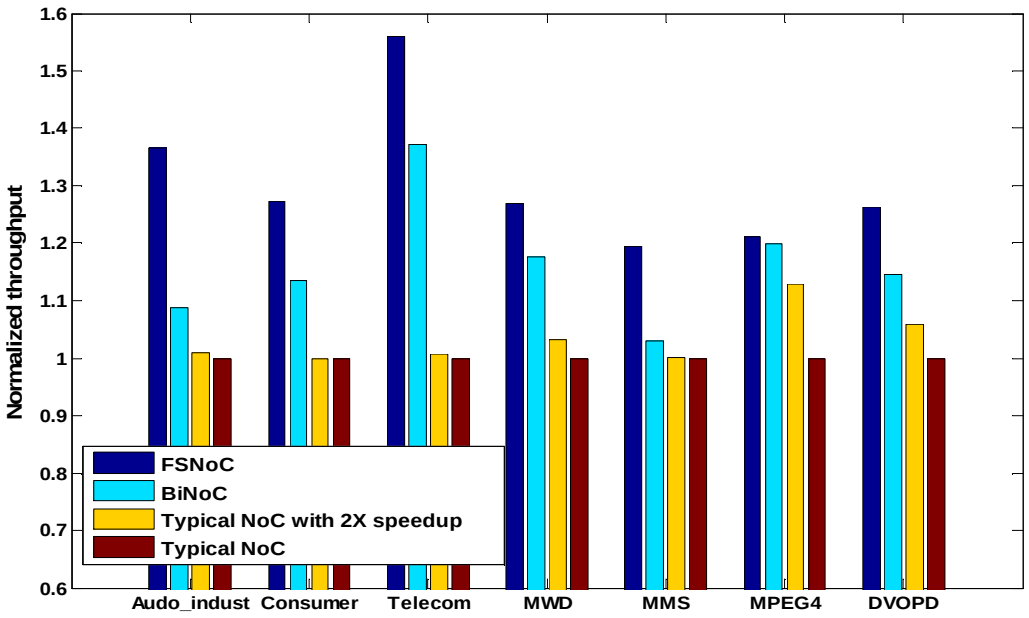

Figure 5.15: Benchmark throughput comparison

traffic patterns such as transpose, the oddeven routing algorithm evenly distribute the load during the routing, therefore, the improvement of FSNoC over BiNoC is smaller than the case of XY routing. However, as can be seen from the figure, the FSNoC still achieves 6.5% improvement in the saturation injection point for the random and transpose traffic. In Fig. 5.14-b, we evaluate the influence of number of virtual channels on the latency performance. For all architectures, the buffer size per input port is fixed as in Table 5.1. As shown in Fig. 5.14-b, FSNoC consistently achieves the highest performance for both the wormhole (WH) architecture and two virtual channel architectures. The idea of flit-level speedup can also be applied to NoCs with four uni-directional links between router pairs (*i.e.,* $4L$ link architecture in [101, 102]). Because the $4L$ link architecture always provides $2X$ bandwidth for each direction, it gives the upper-bound performance that BiNoC can achieve. Conventional $4L$ link design only allows two different packets to use the $2X$ channel. Therefore, we can improve the performance by allowing two flits from the same packet to use $2X$ links (*i.e.,* flit-level speedup). In Fig. 5.14, we compare the latency performance of four architectures, namely BiNoC, FSNoC, Typical NoC with $4L$ link and FSNoC with $4L$ link, using random and transpose traffic patterns. As can be seen from the figure, for uniform traffic patterns such as random, the $4L$ link architectures (*e.g.,* Typical $4L$ and FSNoC $4L$) achieve higher performance than the FSNoC and BiNoC because the static bandwidth assignment avoids the direction switching overhead existed in

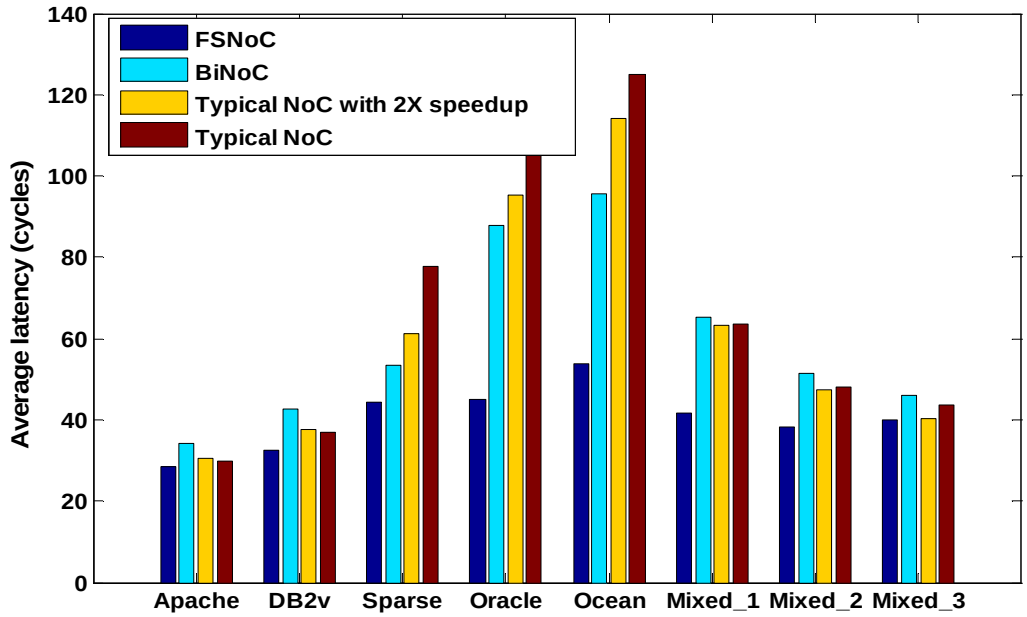

Figure 5.16: Latency comparisons for the SPEC-web benchmarks

bi-directional designs. For uneven traffic patterns such as transpose, the Typical $4L$ and FSNoC $4L$ architectures have similar performance as BiNoC and FSNoC, respectively. This is because the direction switching is not so frequent in these patterns. Hence, the FSNoC and BiNoC architectures work just like the $4L$ architecture under these cases.

5.5.3 Simulation results on real benchmarks

We also evaluate the FSNoC performance using several real benchmarks. In Fig. 5.16, we compare the saturation throughput for seven MPSoC applications. Of note, for these MPSoC benchmarks such as MMS, MWD, we assume the packets are injected according to a Poisson process whose mean rate is scaled by a parameter c according to the communication data volume on the edges of the task graphs. We then tune the scaling factor c to find out the maximum throughput that various architectures can support for the comparison in Fig. 5.16. All the saturation throughputs are normalized to that of the typical NoC design. As shown in the figure, while the bi-directional switching scheme out-performs the two unidirectional schemes, our proposed flit-level speedup scheme consistently further improves the performance. The relative improvement of FSNoC over BiNoC ranges from 2% to 17% depending on the applications. In Fig. 5.16, we use the collected traces from SPEC-web benchmarks to compare the average

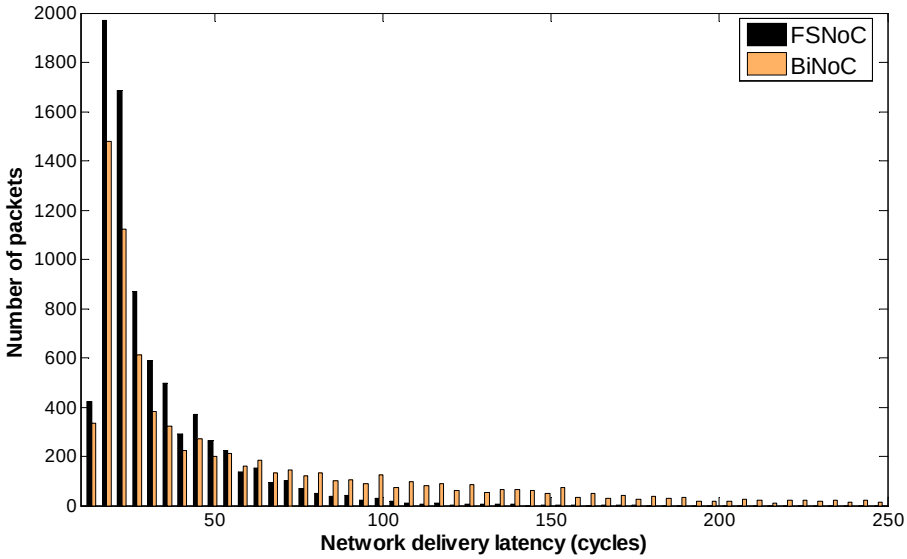

Figure 5.17: Histogram of the delivery time for Telecom

communication latency in the CMP platforms. As shown in the figures, for all the SPEC applications, the proposed FSNoC design achieves fairly good reduction in latency. Especially for those applications with medium or high workload (*e.g.,* *Oracle* and *Ocean* in Fig. 5.16), the improvement is more significant as the FSNoC significantly reduce the network congestion under these workloads by providing more flexibility to use the inter-router bandwidth.

In Figure 5.17, we also show an example histogram of the packet delivery time under 80% saturation injection factor for the telecom application in E3S application. As shown in the figure, not only the average latency but also the maximum delay are significantly reduced. As the worst-case delay is more important for the applications with real-time deadline constraints, the FSNoC also provides a good prospective to satisfy the QoS requirements for these applications.

5.5.4 Simulation results for FSNoC with link pipelines

Regarding the influences of long wire links, we next evaluate the latency performance of four architectures with 2-stage pipeline registers inserted in all links. In Fig. 5.18, the latency simulation results for four synthetic traffic patterns on an 8×8 mesh are shown. As can be seen from the figure, for these two traffic patterns, the network latency has increased at light loads. For example, for the random traffic, comparing Fig. 5.12 with Fig. 5.18, it can be observed the zero

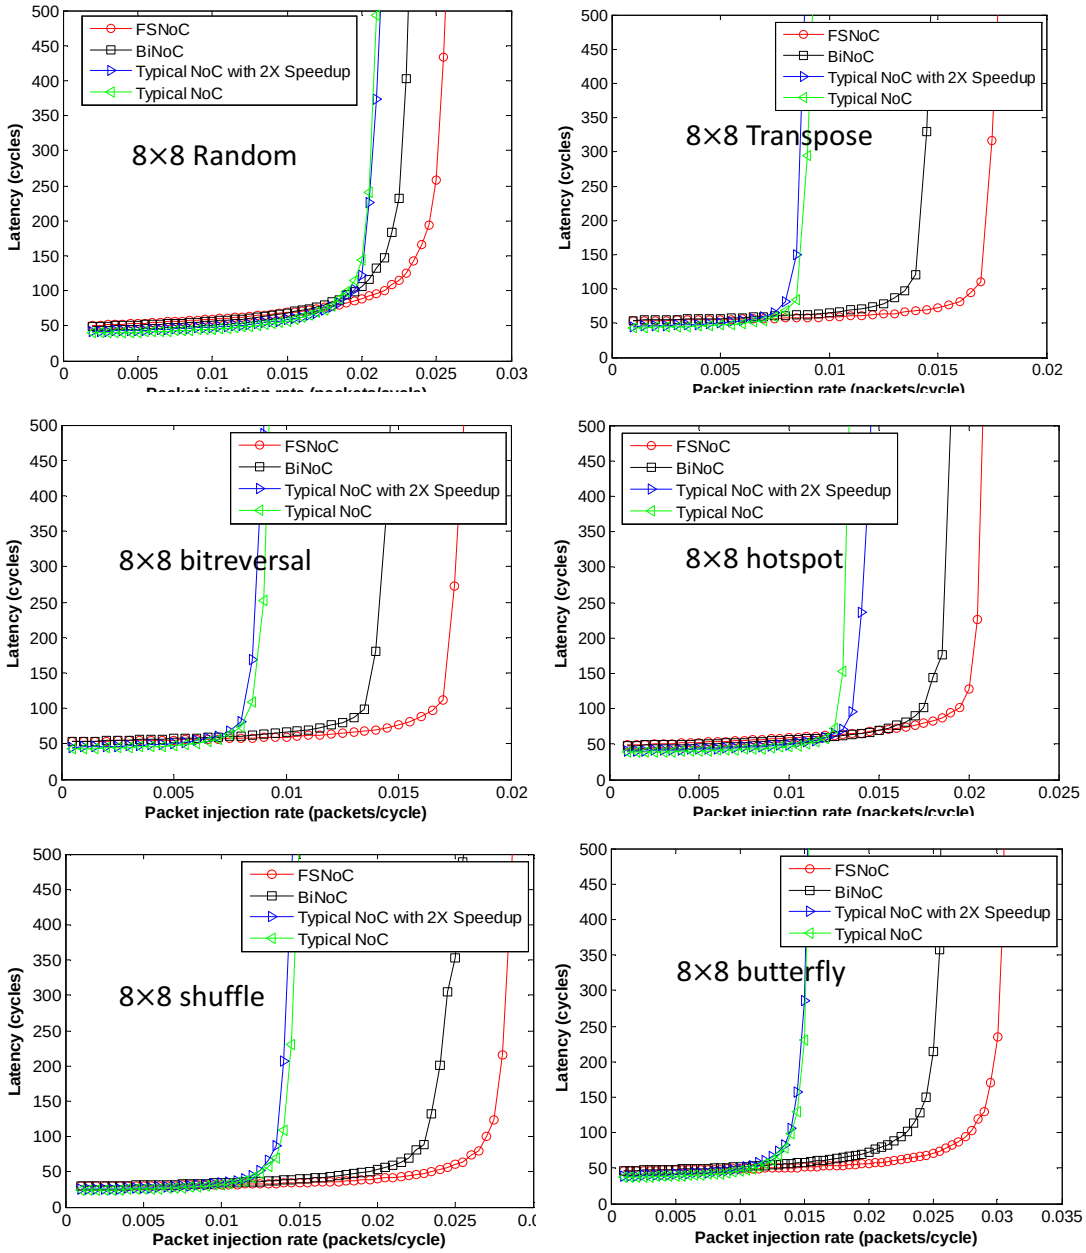

Figure 5.18: Latency comparisons for NoCs with 2-stage pipelined links

load latency has increased from 28 cycles to 40 cycles because it takes 3 cycles in total for a flit to reach the neighboring router. Moreover, the saturation injection point (*i.e., pir*) has also decreased slightly for all the architectures. For random traffic, the *pir* values of four architectures have reduced to 0.018, 0.018, 0.021 and 0.024 *(packet/cycle)*, respectively. This is because in order to avoid both wire and register conflicts described in section 5.4.2, the OWC module should determine the direction control (DS) signals by considering not only the current cycle pressure information but also that of previous two cycles, which reduces the chances to use the $2X$ channel bandwidth. However, as can be seen from the figure, the pipelined FSNoC still outperforms the other three architectures for both random and transpose traffics.

5.5.5 Implementation overhead

As discussed in [101], although the $4L$ link architectures have almost the same overhead as BiNoC in the router design, the more inter-router wiring resources required is not cheap as technologies keep scaling down. Rather than increasing physical bandwidth by adding more additional uni-directional links, NoCs based on bi-directional channels are advocated to reduce the routing congestion and spaces. Therefore, in the hardware overhead evaluation, we majorly implement and compare the four router architectures described in Table 5.1. We implement each scheme in Verilog and synthesized the design using Synopsys Design Vision based on TSMC 65nm library. According to our synthesis results, for all the four architectures, the critical paths are located in the ST stage. For both the BiNoC and FSNoC architectures, the maximum path delay $0.96ns$ is dictated by the 10×10 crossbar and hence they run at the same maximum clock frequency. Table 5.2 and 5.3 summarize the area and power breakdown of different NoC architectures. For FSNoC, the direction control module (i.e., Dir. ctrl. in Table 5.2 and 5.3) includes the Request extractor (RE) and Output width controller (OWC). As shown in the tables, the input buffers consume most of the router area and power for all the four architectures. From the synthesis results, the area overhead of 2X input speedup, BiNoC and the proposed FSNoC over the conventional NoC are 6% , 15% and 20% , respectively. For the power comparison, we assume a 50% switching activity factor of a random payload as in [101]. As can be seen from the table, the power overhead of 2X input speedup, BiNoC and

Table 5.2: Area breakdown of different NoC architectures

	Area(μm^2)(%)							
	Typical NoC		2X speedup		BiNoC		FSNoC	
Input buf.	109774	89.7	110515	88.5	120545	85.2	125268	85.2
RC	1806	1.5	1806	1.4	1806	1.3	1806	1.2
VC allocator	8806	8.0	8806	7.1	8806	6.2	8806	6.0
SA allocator	583	0.5	1138	0.9	1365	1.0	1671	1.1
Crossbar	1429	1.2	2558	2.0	6371	5.3	6371	4.3
Dir. control [a]	*N/A*	*N/A*	*N/A*	*N/A*	2567	1.8	3065	2.1
Total	122398	100	124823	100	141460	100	146989	100
Norm. area	1.0		1.06		1.15		1.20	

Table 5.3: Power breakdown of different NoC architectures

	Power(mw)(%)							
	Typical NoC		2X speedup		BiNoC		FSNoC	
Input buf.	10.30	89.4	10.7	88.4	11.3	87.1	11.5	87.2
RC	0.10	0.9	0.10	0.8	0.10	0.8	0.10	0.8
VC allocator	0.94	8.2	0.94	7.8	0.94	7.2	0.94	7.1
SA allocator	0.04	0.3	0.06	0.5	0.06	0.5	0.06	0.5
Crossbar	0.18	1.6	0.30	2.5	0.52	4.0	0.52	3.9
Dir. control	*N/A*	*N/A*	*N/A*	*N/A*	0.05	0.4	0.06	0.5
Total	11.52	100	12.10	100	12.97	100	13.18	100
Norm. area	1.00		1.05		1.13		1.14	

FSNoC over conventional NoC are 5% , 13% and 14% , respectively.

5.6 Conclusion

In this chapter, we have proposed a flit-level speedup scheme for improving the NoC performance using self-reconfigurable bi-directional links. In order to support transmitting two flits within the same packet at the same cycle, a novel channel direction control protocol is proposed to dynamically configure the link directions. The corresponding design of the input buffer organization and the

switch allocator are also proposed. Also, we have extended the channel direction control scheme to work under the existence of long wire links. From the simulation results, significant improvement in latency and throughput are achieved for both synthetic traffic and the real benchmarks.

CHAPTER 6

A TRAFFIC-AWARE ADAPTIVE ROUTING ALGORITHM ON A HIGHLY RECONFIGURABLE NETWORK-ON-CHIP ARCHITECTURE

In this chapter, we propose a flexible NoC architecture and a dynamic distributed routing algorithm which can enhance the NoC communication performance with minimal energy overhead. In particular, our proposed NoC architecture exploits the following two features: i) self-reconfigurable bidirectional channels to increase the effective bandwidth and ii) express virtual paths, as well as localized hub routers, to bypass some intermediate nodes at run time in the network. A deadlock-free and traffic-aware dynamic routing algorithm is further developed for the proposed architecture, which can take advantage of the increased flexibility in the proposed architecture. Both the channels self-reconfiguration and routing decisions are made in a distributed fashion, based on a function of the localized traffic conditions, in order to maximize the performance and minimize the energy costs at the macroscopic level. Our simulation results show that the proposed approach can reduce the network latency by $30\% - 80\%$ *in most cases compared to a conventional unidirectional mesh topology, while incurring less than 15% power overhead.*

6.1 Introduction

The NoC architecture and the corresponding routing strategy play an important role in optimizing the system performance in terms of throughput and latency. Indeed, enhancing the NoC architecture with additional flexibility is appealing since it can offer more opportunities to minimize the impact of traffic congestion by bypassing the intermediate router pipeline stages for some packets at run time [113]. At the same time, a routing strategy, which fully utilizes the characteristics of both the underlying communication infrastructure and the NoC traffic,

is equally important since the routing decision made at the run time is critical in optimizing the NoC performance and resources utilization [37, 114].

For NoC architecture, the bandwidth and the average distance between the nodes have a significant impact on the overall network performance. Most of the existing NoC architectures have employed a two-dimensional mesh topology with two unidirectional links connecting the neighboring tiles. Although the mesh topology is well suited for silicon implementation [115], the channel bandwidth is not optimally used [102] and it may suffer from long packet latencies due to the lack of short paths between remotely located nodes [108].

For routing strategy design, several adaptive routing algorithms have been proposed for conventional unidirectional mesh NoCs [112, 114–116]. However, dynamic routing has been far less explored for performance improvement.

In this chapter, we propose a reconfigurable NoC architecture which combines the advantages of providing higher effective bandwidth between the neighboring nodes and shorter paths between remotely located nodes [117]. Our proposed NoC architecture exploits the following two capabilities, namely: i) self reconfigurable bidirectional channels (BiNoC) [100, 102] and ii) static express virtual channels (EVCs) [113] with regional hub routers, to dynamically create short-cut paths for transmitting packets. In order to exploit the added flexibility provided by the proposed architecture, we also propose a new fitness-based adaptive routing algorithm for optimizing the network performance.

Experimental results obtained for synthetic traffic and real benchmarks show that our approach can improve the latency by as much as 80% , while involving less than 15% overhead in power dissipation.

The remaining of this chapter is organized as follows: In section 6.2, we review related work on NoC architecture and adaptive routing algorithm design. In section 6.3, we describe the newly proposed NoC architecture. Section 6.4 presents the adaptive routing algorithm for the proposed NoC. The practical considerations and simulation results are discussed in section 6.5. Finally, section 6.6 concludes this work.

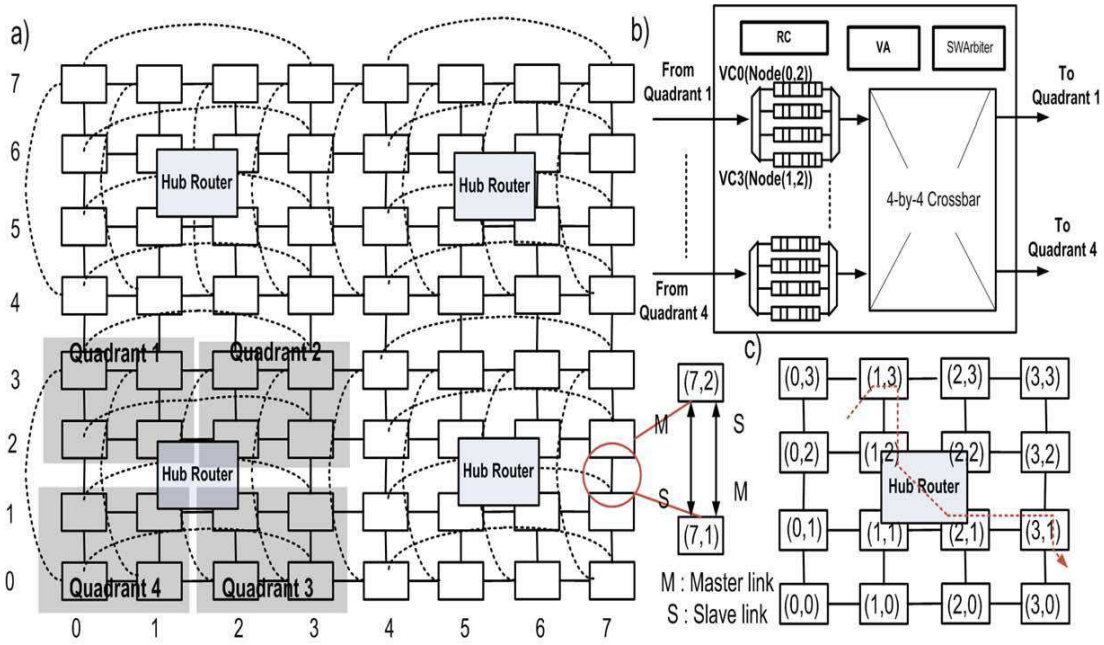

Figure 6.1: a) Proposed NoC architecture for an 8×8 mesh with bi-directional links, EVCs and hub routers b) The hub router architecture with each VC corresponding to a node in the region c) An example of routing from (1,3) to (3,1) using hub routers

6.2 Background

In NoC architecture design, many efforts are focused on increasing the total bandwidth and reducing the average node distance by proposing new topologies [118,119]. In [118], a Diagonally-linked Mesh topology is proposed which employs physical diagonal express links between routers to reduce the distances between nodes. In [119], a new topology called Multidrop Express Channels (MECS) is proposed; this uses a one-to-many communication model to enable a high degree of connectivity. Compared to the mesh based NoC, these new topologies can improve the network bandwidth and connectivity at the cost of adding more physical resources such as the link wires and increased the router complexity.

To improve the trade-off between the design complexity and the performance, there are also approaches exploiting the mesh topology due to its modularity and scalability. In [113], the NoC architecture with static and dynamic express channels (EVC) is proposed. Instead of adding physical links, a flow control mechanism allows packets to virtually bypass intermediate routers along their path. Static EVC approach uses the express paths of uniform lengths and distinguishes the nodes as either EVC source (sink) nodes or bypass nodes according to

whether the EVC originates (terminates) at those nodes. Dynamic EVCs make every node in the network a source (sink) and allow EVCs of various lengths to originate from a node [113].

Besides the idea of bypassing intermediate routers, another important direction in optimizing the mesh NoC architecture is to enable the network to increase the effective bandwidth at run time. Indeed, very often, there is much traffic in one direction, while the channel for the opposite direction is idle. NoCs with run-time reconfigurable bi-directional channels have been recently proposed to fully utilize the bandwidth according to the run-time traffic conditions (BiNoC) [100–102,120]. However, adaptive routing is less exploited in these works for simplicity reasons.

Previous adaptive routing algorithms are mainly designed for unidirectional NoCs. In [114], a neighbors-on-path (NoP) selection strategy is proposed to make each node routing selection based on the condition of the nodes adjacent to the neighbors. In [115], a regional contention awareness scheme (RCA) was first proposed to utilize both the local and non-local information to improve the load balancing in NoCs. In [112], the RCA scheme was improved by using the adaptive routing scheme DBAR to leverage the local and non-local network information. In [116], a destination-based adaptive routing algorithm (DAR) is proposed where every node estimates the delay to every other node in the network. However, for a more flexible NoC platform such as the ones employing bi-directional links and EVC channels, these schemes need to be modified with a congestion fitness function to reflect the run time traffic status.

Starting from these overarching ideas, in this work, we aim at improving the NoC performance with respect to both architecture design and routing algorithm development. Towards this end, we propose a flexible NoC architecture which combines the advantages of higher effective bandwidth and router bypass capability and develop an adaptive routing algorithm to optimize the network run-time performance.

6.3 New NoC architecture

In this section, we first present the motivation for adding regional hub routers into the NoC and then elaborate on the new reconfigurable NoC architecture.

6.3.1 Motivation for adding hub routers

Because of the complexity of the dynamic EVC approach [113], in this work, we adopt a 3-hop static EVC network; this provides a better trade off between performance and design complexity. We also define a region as a 4×4 mesh whose boundary is made up of the EVC source/sink nodes. A regional hub router is added into each region based on the following key observations:

1) The regional traffic represents a significant portion of the total traffic distribution. Here, we define the regional traffic as the communication flow whose source and destination reside within the same region. This is especially true for application-specific NoCs where highly communicating nodes are usually mapped close to each other. For example, in [101], the regional traffic with more than 90% of nodes communicating with other nodes residing within a region of three hops away is used in the evaluation.

2) While the static EVCs provide express paths crossing over NoC regions, they typically lack sufficient paths needed to handle the traffic within the NoC regions. If both the source and destination nodes are bypass nodes in a static EVC, then the packets cannot use the EVCs to bypass the intermediate nodes.

3) The static EVCs enable packets to bypass the intermediate nodes following a fixed express path. Also, the EVCs are restricted to a single dimension and cannot turn directly. Consequently, it will be beneficial if we can dynamically build shortcut paths for regional traffic rather than being limited by the single dimension requirement.

Therefore, in this work, we propose to add a light weight hub router to minimize the burden of the regional traffic in NoCs. More precisely, we additionally allocate a dedicated (*i.e.,* static) express path for each node to route to the hub router and utilize the hub router to connect the nodes within the same region; this way, when the hub router is used, the intermediate pipeline stages of the normal routers can be bypassed altogether.

6.3.2 Overview of proposed NoC architecture

Figure 6.1-a shows the proposed flexible NoC architecture for 8×8 meshes. In this example, the static EVC hop length is equal to 3. Each region is a small 4×4 mesh. Other static EVC hop length and region size can be used for different sizes of the meshes. The solid line in Figure 6.1-a represents the physical channels between the tiles which are made up of two bi-directional links. Both bidirectional links can be reconfigured as sending or receiving. For each router, we define one link as a high priority link (master link) and the other as a low priority link (slave link) as in [101]. By default, the master link is the sending link and the slave link is the receiving link when there is traffic in both directions. As shown in Figure 6.1-a, the master link of one router corresponds to the slave link of its neighbor. The dash line in the figure represents the 3-hop EVCs.

We further divide each 4×4 region into four quadrants (see Figure 6.1-a). In order to reduce the additional complexity added to the NoC backbone, the regional hub router is a simple and generic four ports router with four virtual channels at each input port as shown in Figure 6.1-b. Every input port corresponds to a quadrant in the region and each node in the quadrant is assigned a virtual channel at the input port of the hub router. For example, VC0 of the first input port (i.e. the quadrant 1 port) of the hub router is assigned to node (0,2) in Figure 6.1-a.

At run time, the hub router plays the role of setting a connection between the source and destination pairs within different quadrants of the same region. For example, in Figure 6.1-c, if node (1,3) in quadrant one needs to communicate with node (3,1) in the quadrant three, it cannot use the static EVCs. Instead, it can communicate via the hub router. The packet will first follow the dedicated express hub path bypassing router (1,2) pipeline stages to reach the input port corresponding to quadrant one of the hub router. After the hub router grants its connection from quadrant one input port to the quadrant three output port in the switch allocation stage, the packet will traverse the crossbar in the hub router and bypass router (2,1) pipeline stages to arrive at the destination node (3,1) directly. Compared with using normal paths to send packets from (1,3) to (3,1), which requires five router pipeline stages, only one router pipeline in the hub router is needed.

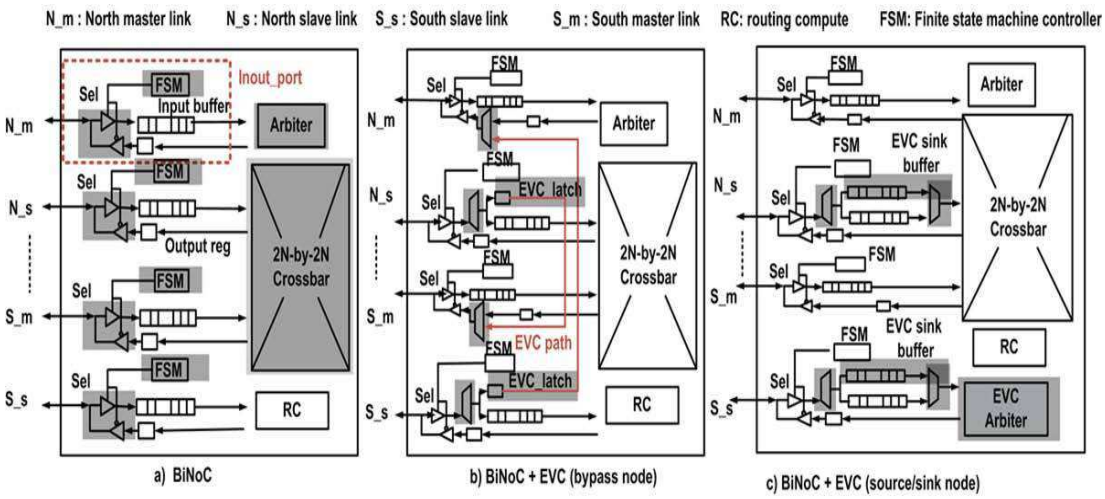

Figure 6.2: Router architecture of a) BiNoC; b) BiNoC + EVC bypass node; c) BiNoC+EVC source/sink node

In the proposed architecture, we utilize the virtual channel allocator and switch allocator in the hub router to arbitrate the contentions. For instance, if both nodes (1,3) and (1,0) want to send packets to (3,1) using the hub router, after arriving at the hub router, these two packets will request the same virtual channel which corresponds to node (3,1) in the quadrant three output port. The virtual channel allocator will grant one of the requests first, while the other packet needs to wait in the input channel buffer until the virtual channel becomes available again. Similarly, when there are multiple requests from different input quadrant ports to the same output quadrant, for instance, if both the packets from node (1,3) to (3,1) and node (1,1) to (3,0) want to use the quadrant four output port in the hub router, the switch allocator will only serve one request each time.

6.3.3 Router implementation

In this section, we present the router design that supports the BiNoC with EVC and region hub. First, we will discuss the additional hardware required to support the BiNoC approach. Then the micro-architecture that supports static EVCs will be presented. Finally, we present the modifications required to support the region-hub routing.

Figure 6.2-a shows the router design for BiNoC with the extra logic required (see the shaded area) to support the bi-directional switching. Compared to a unidirectional wormhole router, in BiNoC each input and output port need to

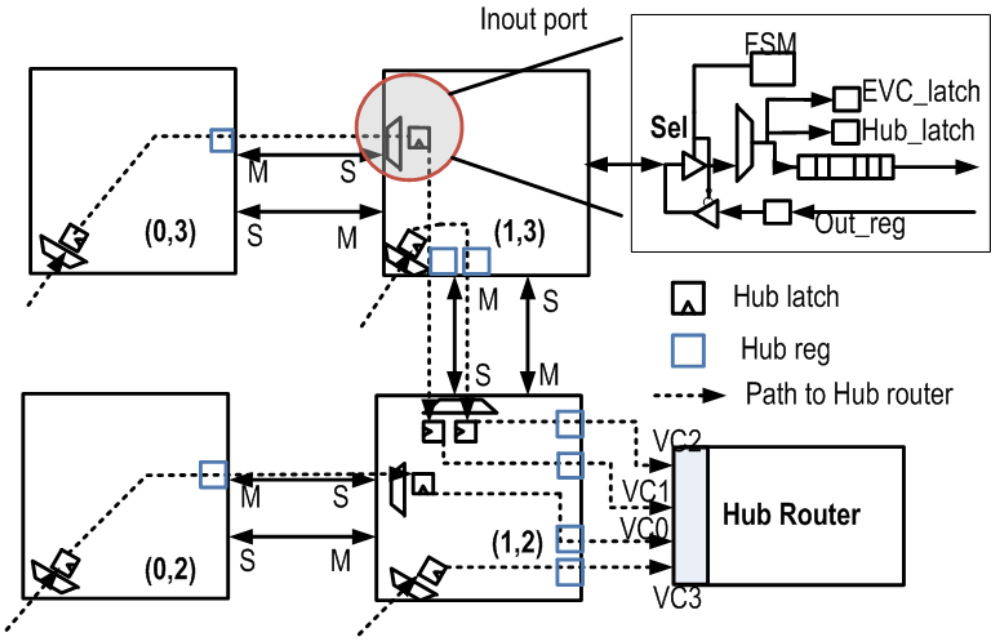

Figure 6.3: The express paths to the hub router

be modified to inout ports controlled by the corresponding finite state machines (FSMs) [101]. Also, the arbiter and the crossbar are modified to handle the requests received from the master and slave ports. At run time, the FSM associated to every inout port configures the port mode (*i.e.,* input or output) according to the requests existing in current router and the backpressure feedback from the neighboring routers as discussed in [101].

Figure 6.2-b shows the router micro-architecture that supports static EVCs for BiNoC. In this work, we build the EVC path using the output ports corresponding to the master links, because the master link has a higher priority in transmitting flits according to the channel direction controlled by the FSM. Compared to the BiNoC router in Figure 6.2-a, the extra logic need to support the EVC architecture is shaded in Figure 6.2-b and Figure 6.2-c. In static EVCs, there are two types of nodes, namely the bypass nodes and the EVC source/sink nodes. In the bypass nodes (as shown in Figure 6.2-b), an EVC latch is added to hold the packets using the EVC. Once a packet arrives at the EVC latch, the packet will bypass all the intermediate pipeline stages of the current router (such as switch allocation, crossbar traversal) and is directly sent to the output port in the same dimension. In Figure 6.2-b, the EVC paths along the north/south direction that a packet will take to bypass the router pipeline is highlighted. The architecture of the EVC source/sink routers is shown in Figure 6.2-c. Here, we add an EVC

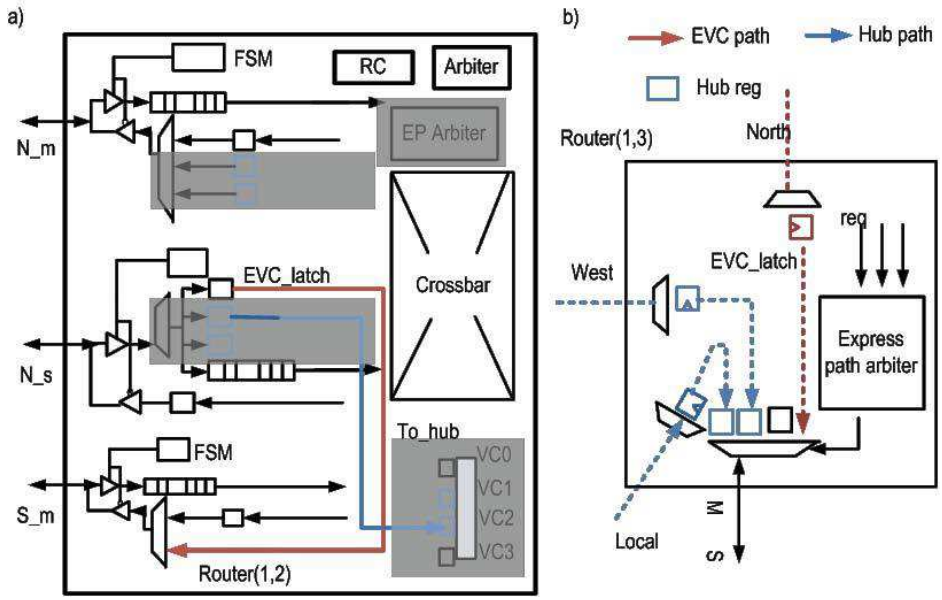

Figure 6.4: a)Router architecture of node (1,2) b) Contention between EVC and Hub paths

sink buffer at the slave link port to hold the packets from the upstream nodes. In addition, an EVC arbiter is added to handle the arbitration requests from all the input channels (both the normal input and EVC sink buffers). Once a packet successfully wins the allocation, an EVC flag in the flit will be asserted. Based on this flag, the downstream router will decide whether to put the flits into the EVC latch or into the normal input buffer.

In the proposed regional hub routing, for every node, there is one static express path assigned to connect to the hub router. Figure 6.3 shows the paths assigned to the four nodes in the quadrant one of Figure 6.1. We add hub latches to multiplex with the normal buffers at the input port and hub registers to multiplex with the output registers at the output port (shown in Figure 6.3). The hub latch and register pair forms an express path for a dedicated node to bypass the pipeline stage within the router.

Figure 6.4-a shows the router design for node (1,2) in Figure 6.3 which connects to the hub router using an additional port. The extra logic needed compared to Figure 6.2-c is shaded. As shown in Figure 6.4-a, the EVC paths are highlighted in red lines while the express paths formed by hub latch/register pairs are highlighted in blue lines.

At run time, there may be contentions between the hub latches or between the EVC and hub latches. For example, for router (1,3) in Figure 6.4-b, the hub

latch in the local port and the hub latch in the west port may contend with the EVC latch in the north port to use the south output link. In order to resolve such contentions, the express path arbiter in Figure 6.4-b is needed in every router to handle the hub and EVC latch requests to make sure that, at any one time, only one packet can bypass the pipeline stages and use the physical channel. In this work, we allocate a higher priority to the EVC paths, because the hub paths only serve the regional traffic, while the EVC paths usually work for the traffic meant for longer distances. For example, the express path arbiter will first grant the north EVC latch to use the south output link. The packets from the west and local hub latches, as well as the input buffers will wait in the output port. Then, the express path arbiter will grant the requests from hub latches to use the output link with a higher priority than the normal paths. In order to avoid the starvation of flits in normal buffers and registers, the arbiter will grant the requests from normal paths after serving the EVC and hub latches for n consecutive cycles as in [113]. In this work, n is set to be 30 according to our simulation results.

6.4 Routing algorithm design

In this section, we present the details of our adaptive routing algorithm for the proposed flexible NoC. We first describe the scheme to choose among the static EVCs, express hub paths and BiNoC for sending a packet. Next, we present the adaptive routing algorithm when using the bi-directional channels to route packets. We develop a fitness function as a metric to evaluate the neighboring node availability based on the channels occupancy. The fitness function not only captures the buffer status but also takes the traffic dynamics, as well as the characteristics of the bi-directional channel into account. Then, a K-step speculative routing selection strategy is proposed to evaluate each candidate direction, which dynamically selects the nodes with the best fitness values that reside on a path towards the destination.

6.4.1 Choosing the express hub and EVC paths

In the proposed NoC architecture, we can use EVCs, hub routers, as well as the bi-directional channels to send packets around. The network interface (NI) in the source tile determines whether to use the express hub paths for the packet first.

More specifically, the NI first compares the source and destination address of the packet. If the source and destination nodes are located in different quadrants of the same region, and if the destination node cannot be reached by EVCs from the source directly, then the NI will check the availability of the hub latch in the local port of the neighboring router. If the hub latch is free, the packet will be sent to it and follow the express hub path to route towards the destination. Otherwise, the packet will be sent to the normal buffer in the local port of the neighboring router.

In the EVC source/sink nodes, for each input channel, the routing computation module decides whether to use the EVCs based on the EVC latch availability and the average waiting time of the sink buffer in the EVC sink node. If the waiting time is larger than a threshold T (which means that there exists heavy congestion at the EVC sink node), the EVC will not be chosen for routing.

At last, for the packets which cannot use EVCs or the express hub paths, the adaptive routing algorithm proposed for bi-directional NoC is used to find an output direction based on the run time traffic status.

6.4.2 Adaptive routing for BiNoC

Channel and direction fitness

In NoCs utilizing bidirectional channels, for each direction, there are two input/output channels corresponding to the master and slave links, respectively. We use two metrics, namely channel fitness and direction fitness, to reflect the suitability of an input channel to receive packets.

In the channel fitness function, in order to predict the dynamics of the network traffic which may exhibit a non-Markovian behavior [37], we include the channel average waiting time to reflect the channels "memory". The fitness value of an input channel is represented as:

$$fitness_{in} = \alpha \times m_{in} \times w_{in}^{-1} \tag{6.1}$$

where α is a scaling parameter, m_{in} is the average number of free slots in the channel and w_{in} is the average waiting time that a packet spends in the channel. In this work, according to the simulation results with different α values ranging from 0.5 to 1.5, we observe α equals to 0.8 offers the highest performance.

Table 6.1: Notations in the routing for proposed NoC

Parameter	Description
N	Number of nodes in the mesh network
P	The set of directions $P = N, E, S, W, L$
$f[i][j][k]$	Fitness of the input channel where $i \in P$, $j \in J$, $k \in N$
$P_{block}[i][j]$	Blocking probability of the input channel in master inout port in direction i of node k
$F[i][k]$	Fitness value of direction i of node k
$S_l[i][k]$	Total score from current node l to node k through direction i
$CH[i][k]$	Set of (node, direction) pairs chosen by K-step speculative algorithm from current node to the destination node k along direction i

During the routing stage, since we have not yet determined whether the master or the slave output port will be used, we need to evaluate the input channels at both ports of the downstream router to get a more appropriate metric for choosing the output direction. The direction fitness function is used to indicate the overall availability of one direction in the downstream router. One issue when designing the direction fitness function is to consider the characteristics of the bi-directional channels since the effective bandwidth between two neighboring routers varies with the actual traffic over time. One simple way is to add together the two channels fitness values regardless of them being of master or slave type. However, because the master port has a higher priority in sending and its input channel can only receive flits when there is no sending request at the same port. Then, if we only consider the channel fitness without regarding to the probability of receiving flits using this channel or the actual channel utilization, it may not lead to an optimal selection.

In order to deal with this issue, we calculate the fitness of a particular direction by distinguishing the input channels at the master and slave ports. For the input channel in the slave port, because the channel can always be used whenever there is a request from a neighbor, its fitness directly contributes to the fitness function calculation. For the input channels at the master inout port, we calculate the average block probability P_{block} due to the write conflict during a time window. This is the probability when the FSM controller configures the port into the

Algorithm 5 Channel direction fitness calculation

Input: Average waiting time: $w[i][j][k]$; Average available slots $m[i][j][k]$; master port input channel block probability $P_{block}[i][j]$; scaling parameter α
Output: Channel fitness: $f[i][j][k]$ and the direction fitness function $F[i][k]$

1: **for all** node $k \in N$ **do**
2: **for all** $i \in P$ **do**
3: $F[i][k] = 0$
4: **for all** $j \in P$ **do**
4: $f[i][j][k] = \alpha \times m[i][j][k] \times w[i][j][k]^{-1}$
5: **if** $j == 0$ **then**
6: $F[i][k] += (1 - P_{block}[i][j]) \times f[i][j][k]$
7: **else**
8: $F[i][k] += f[i][j][k]$
9: **end if**
10: **end for**{for all input direction i}
11: **end for**{for all output direction j}
12: **end for**{for all the nodes in the NoC}

output mode and lets the output register use the link for sending packets instead of receiving the packets from the neighboring nodes. The fitness of direction i under consideration is then:

$$F_i = fitness_{slave\_in} + (1 - P_{block}) \times fitness_{master\_in} \tag{6.2}$$

In Algorithm 5, we present the idea of calculating the fitness function periodically for any given direction. The notations used for the algorithm are listed in Table 6.1. Two counters are needed in each input channel to calculate the average waiting time of packets and the number of available buffer slots during any time window. According to these two metrics, the input channel fitness at the master and slave inout ports can be calculated. An extra counter at the master port FSM is used to record the average blocking probability P_{block}. The fitness of the direction under evaluation is then the weighted sum of the two input channels fitness values.

Fitness based adaptive routing algorithm

In order to avoid deadlocks at run time, the odd-even turn rules [91] are used in the minimal path based routing computation. Under this assumption, for any packet in the current router, there are at most two feasible output directions.

As discussed in [114–116], if the adaptive routing algorithm only considers the local congestion metric, it will be too slow to re-act to congestion in the more distant parts of the network as it relies on network backpressure to propagate

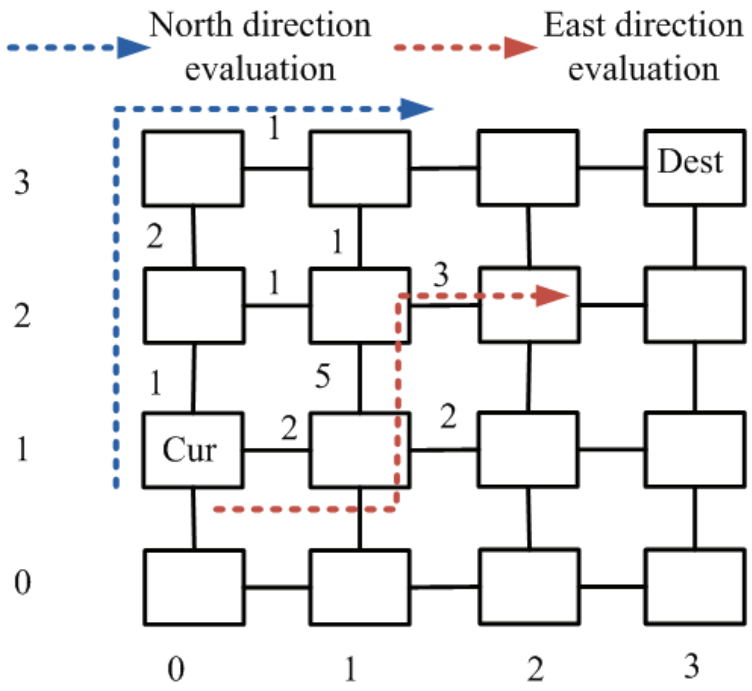

Figure 6.5: K-step (K=3) speculative routing

the congestion state. In this work, we propose a scheme that combines the local and remote channel fitness status in order to make a better routing decision (see Algorithm 2).

Different from previous approaches in [112, 114, 115], which utilize fixed local and non-local neighbors to evaluate each candidate direction in the router, our proposed adaptive routing algorithm uses a K-step speculative routing to dynamically select a set of nodes on the routing path which have the best fitness value towards the destination currently. Based on the fitness value of these nodes, the routing algorithm decides the optimal output port for the current router.

Figure 6.5 shows a routing scenario for the proposed K-step speculative routing with $K = 3$. The values on the link indicate the direction fitness level measured at each router. When evaluating the east output direction for the current node, node (1,1) is chosen first. Then, the router compares the two candidate directions from node (1,1) to the destination. Since the north output direction at node (1,1) has a larger direction fitness value (5) than the east output direction. The north neighboring node (1,2) is chosen second. Similarly, for node (1,2), its east neighboring node (2,2) will be selected. Therefore, nodes (1,1), (1,2) and (2,2) form a 3-step speculative path towards the destination using the east output direction of the current router. Similarly, nodes (0,2), (0,3) and (1,3) form a

speculative path from the north output direction towards the destination.

Algorithm 6 K-step speculative selection algorithm

Input: Current node **cur**: Destination node **dst**; Set of candidate output directions P_c
Output: Output direction $P_o \in P_c$
Function:$Candidate\_dir(a, b)$: A set of candidate directions for node a

```
1: for all direction i ∈ P_c do
2:    S_cur[i][dst] = 0;
3:    CH[i][dst] ← (cur, i);
4:    check_length = 0;
5:    while check_length < K&&!CH[i][dst].IsEmpty() do
6:       (check_node, check_dir) = CH[i][dst].Pop();
7:       next_node = check_node.get_neighbor(check_dir);
8:       if next_node! = dst then
9:          S_cur[i][dst]+ = F[reverse(check_dir)][next_node];
10:         C_set = Candidate_dir(next_node, dst);
11:         if F[C_set[1]][next_node] > F[C_set[0]][next_node] then
12:            CH[i][dst] ← (next_node, C_set[1]);
13:         else
14:            CH[i][dst] ← (next_node, C_set[0]);
15:         end if
16:      end if
17:      check_length + +;
18:   end while
19: end for
20: P_0 = Select_highest_score_dir(S_cur, P_c);
```

By adding together the direction fitness values of each path, the scores of using the north and the east output port can be calculated. The router then chooses the direction with the best score for the current router. If the direction fitness value changes at run time, the nodes selected for evaluating each candidate port will also change accordingly to capture the network dynamics.

Algorithm 6 presents the details of the proposed K-step speculative routing algorithm. At run time, each router periodically collects the channel and direction fitness values of the neighbors located in a region less than K-hops away. Then, for each output candidate, we perform a K-step look ahead routing towards the destinations and select a set of K consecutive nodes and directions to form a speculative routing path towards the destinations using current candidate direction. By comparing the total fitness values along the paths between the two candidate directions, the direction with the best path fitness value is chosen as the output port for the current router.

In this work, we adopt a light weight traffic status propagation network similar

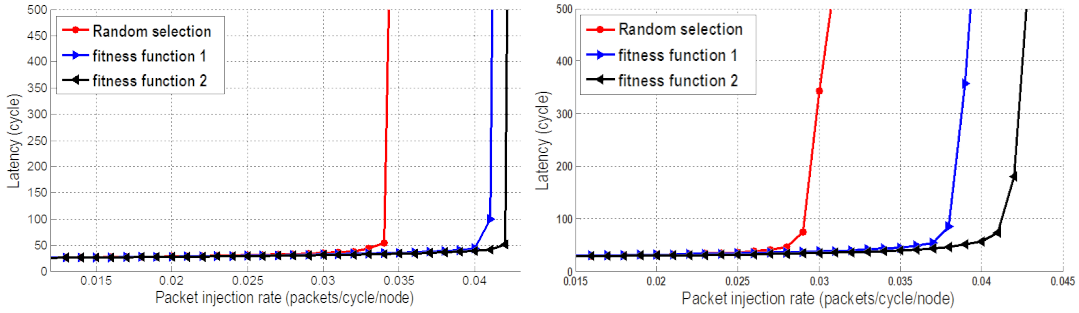

Figure 6.6: Comparison of the fitness functions for 8×8 meshes using bi-directional channels

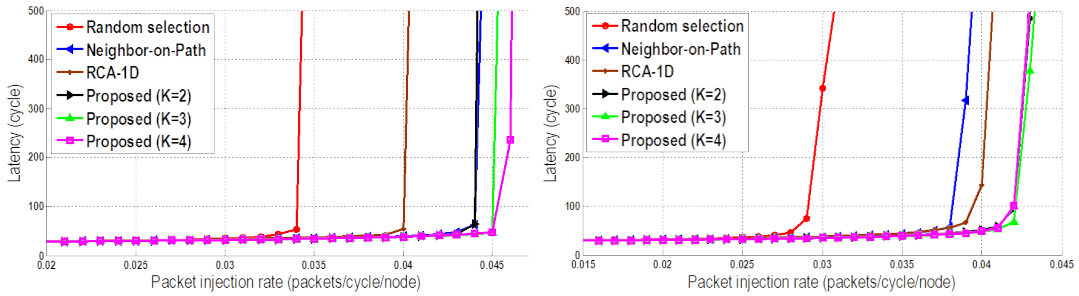

Figure 6.7: Comparison of different adaptive routing schemes for 8×8 meshes

to [116] which is used to propagate the direction fitness information to the nodes within K-hops distance in the network. More specifically, every T cycles, the direction fitness value updated by the current router is propagated K hops away throughout the network in a hop-by-hop manner. After collecting all the direction fitness values from the neighboring routers within K-hops distance, these values are stored in the router and used for the proposed K-step speculative routing algorithm. For the case $K \leq 2$, we can further reduce the network complexity by using a specific control channel between the neighboring nodes as in [114] instead of a status propagation network to exchange the traffic information. The signal sent through the control channel contains the fitness of the direction $F[i][k]$ associated to the current router and its 1-hop neighbors. According to the simulation results shown in section 5, there is not too much performance improvement for $K \geq 2$ compared to $K = 2$. Hence we propose to set $K = 2$ (i.e. look only two hops ahead) to achieve the best trade-off between the performance and additional complexity.

6.5 Simulation results

6.5.1 Simulation setup

We evaluate the proposed NoC architecture, as well as the adaptive routing algorithm, using a cycle accurate NoC simulator implemented in C++. Simulations were done for different mesh network sizes: 4×4 and 8×8. To investigate the benefits of our NoC architecture and the performance improvements of our adaptive traffic aware routing algorithm, we compare our approach with the following NoC architectures: 1) classical unidirectional NoC; 2) bi-directional NoC [101]; 3) bi-directional NoC with static EVCs [113].

Throughout the experiments, we assume that each input channel has a buffer depth of 6 flits for the NoC using bi-directional channels and a buffer depth of 12 flits for the uni-directional NoC. Both synthetic traffic (*i.e.,* uniform, transpose) and real benchmark traces were used in the simulations. For uniform traffic, each node sends packets to other nodes in the network with an equal probability. For transpose traffic, each node at the (i, j) mesh location sends packets to the node at the (j, i) location on the mesh.

For the real benchmark workloads, we consider several 16-node multithreaded commercial and scientific workloads from DBmbench [121] and SPECweb99 [111]. This type of traffic includes the memory request/response coherence traffic between PEs and caches.

6.5.2 Comparison of adaptive routing for BiNoC

We compare our proposed adaptive routing algorithm with some widely adopted routing algorithms to evaluate the performance under bi-directional channels. More precisely, we compare the benefits of using a fitness function and a K-step speculative selection algorithm to capture the network traffic status. First, we compare different fitness functions used to evaluate the router channels. In this comparison, we only gather the traffic status from the immediate neighbors and use different fitness functions to evaluate the channels status. The baseline for this comparison is the random selection strategy of the Odd-even routing algorithm [91]. Two direction fitness functions are used in the simulation:

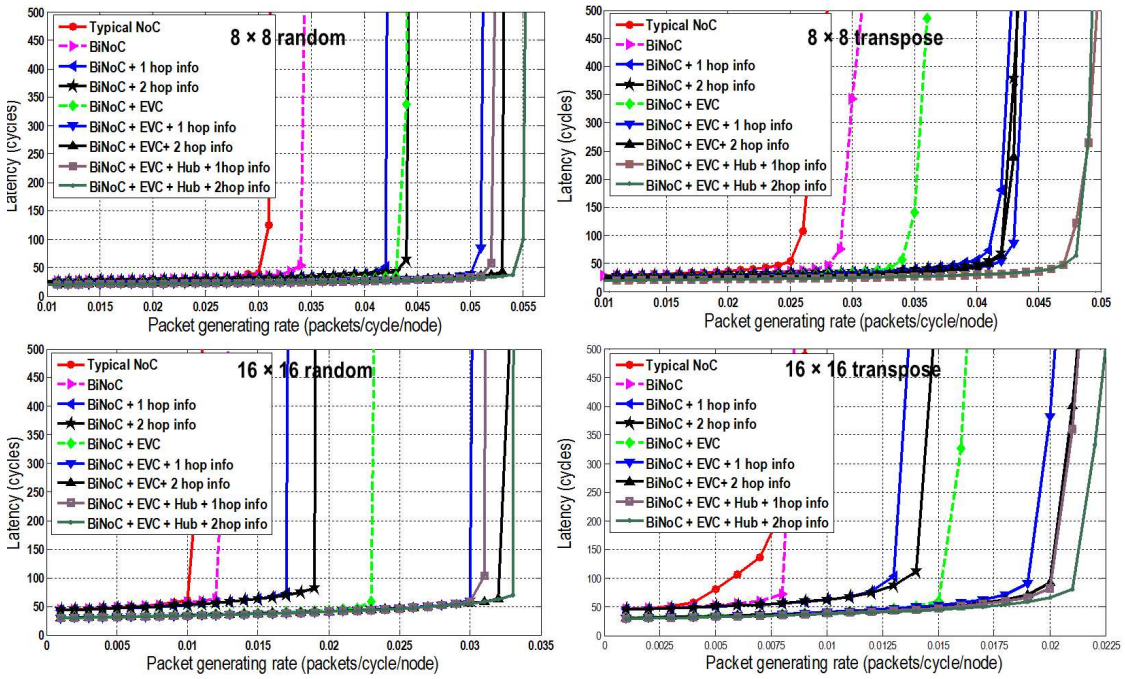
8 × 8 random
8 × 8 transpose
16 × 16 random
16 × 16 transpose
Typical NoC
BiNoC
BiNoC + 1 hop info
BiNoC + 2 hop info
BiNoC + EVC
BiNoC + EVC + 1 hop info
BiNoC + EVC+ 2 hop info
BiNoC + EVC + Hub + 1hop info
BiNoC + EVC + Hub + 2hop info
Latency (cycles)
Packet generating rate (packets/cycle/node)

Figure 6.8: Comparison of different NoC architectures under various mesh size and traffic patterns.

1) Fitness function 1 is the direct sum of the input channel fitness without distinguishing the master and slave channels

2) Fitness function 2 uses the link blocking probability to weight the channels in the combination of the master and slave channels fitness values.

Figure 6.6 shows the comparison results obtained for an 8×8 mesh under random and transpose traffic patterns. As shown in this figure, using a fitness function to represent the channel traffic status significantly improves the network critical load (by more than 22.5%). Furthermore, if we consider the characteristic of the bidirectional links by distinguishing the master and slave channels, an extra 5% - 8% performance gain can be achieved.

Next, using the proposed fitness function for the bidirectional NoC, we make a comparison of different adaptive routing schemes which consider both the immediate neighbors and non-local nodes to evaluate the directions in the routing selection phase. We evaluate our proposed K-step speculative selection scheme with different K values (K=2, 3, 4). At the same time we compare our scheme with two widely adopted schemes, namely neighbor-on-path [114] and regional-contention-awareness (RCA-1D) [115] under random, and transpose traffic patterns on an 8×8 mesh. In order to minimize the intra-region interference in

RCA-1D [112, 116], we modify the RCA-1D scheme to integrate the destination into the selection function as in the DBAR approach [112].

As shown from Figure 6.7, the proposed K-step speculative scheme achieves the best performance results by selectively choosing the intermediate nodes towards destination at run time. Also, it can be seen that there is no significant improvement in the performance for a K value larger than 2. Since the complexity of the fitness value propagation increases dramatically when $K \geq 2$, we use $K = 2$ throughout the remaining simulations.

6.5.3 Effect of adding EVC and hub routers

In this section, we evaluate the performance improvement resulted by adding EVC and regional hub routers into the BiNoC. Both random selection and K-step speculative selection strategies are used in these simulations. The following architectures were evaluated and compared: the typical NoC with unidirectional channels, the BiNoC proposed in [101], the BiNoC + EVC which combines the features of Bi-directional switching in [101] and express virtual channels in [113], and the proposed NoC with regional hubs.

Figure 6.8 summarizes the simulation results. We compare the average latency performance under various traffic injection rate for 4×4 and 8×8 meshes. Several observations can be made based on these results:

1) BiNoC outperforms the typical NoC design. Adding EVC into the network can further improve the network criticality by 9.8%-80% for different traffic patterns and mesh sizes. Furthermore, if the regional hub router is added, the throughput can be further improved by 9.3%-15%.

2) For all the NoC architectures using the adaptive routing algorithm, the fitness function employed to evaluate the traffic status can significantly improve the network performance. As shown in Figure 6.8, using the 1 hop away neighboring info can improve the network criticality by 13% - 30%, while using 2 hop neighbor info can further improve the performance by 3% - 11% in most cases.

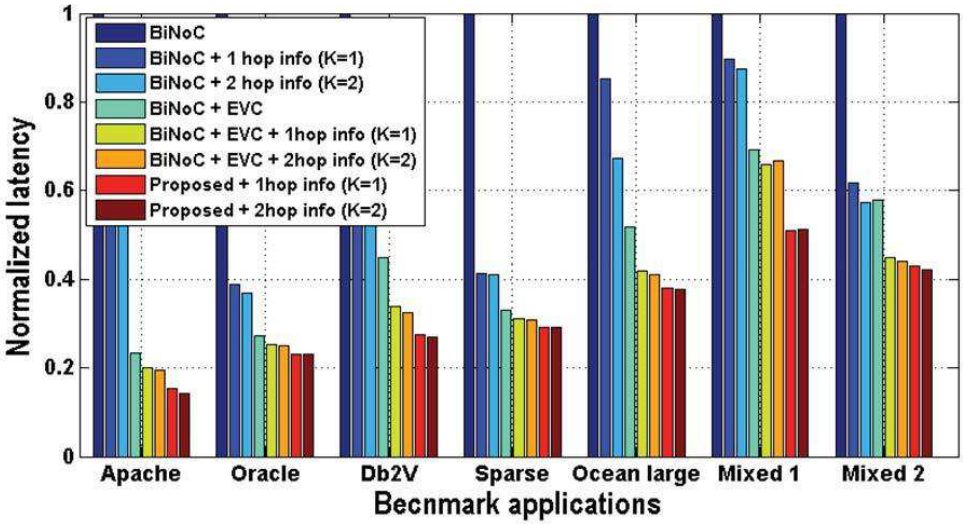

Figure 6.9: Latency comparison for benchmarks

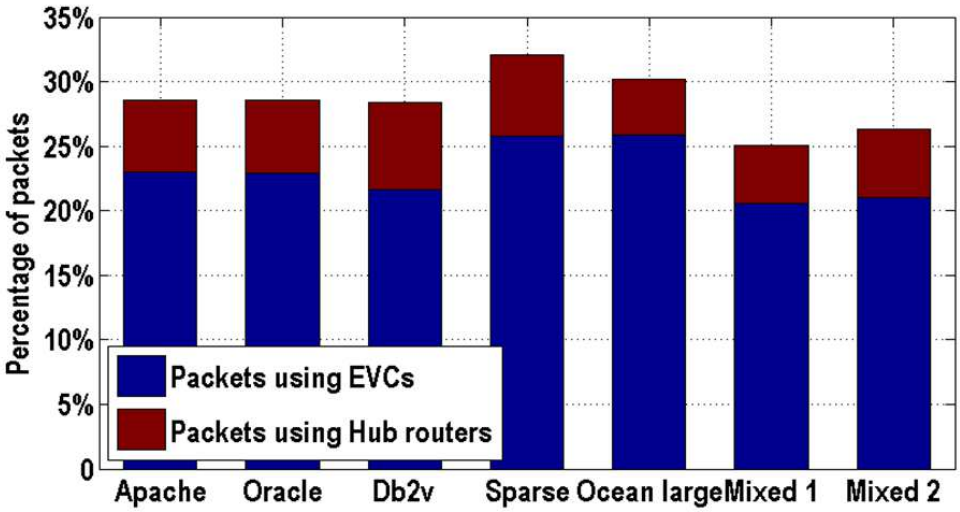

Figure 6.10: Percentage of packets using EVCs and hub routers

6.5.4 Results using real world workloads

In this section, we show the latency comparison for various real world benchmarks with an 8×8 mesh size. We report seven experiments. The first five experiments collect the traces from the same applications, *i.e.,* Apache, Oracle, Db2V, sparse and Ocean which were mapped onto an 8×8 mesh. The last two experiments named Mixed-1 and Mixed-2 combine the traces from different applications and map them to the same NoC platform. The simulation time of each benchmark is set to $5,000,000$ clock circles with a $2,000$ cycles warmup period.

As shown in Figure 6.9, compared with the baseline BiNoC design with random selection strategy, which is normalized to 1, by adopting a fitness function

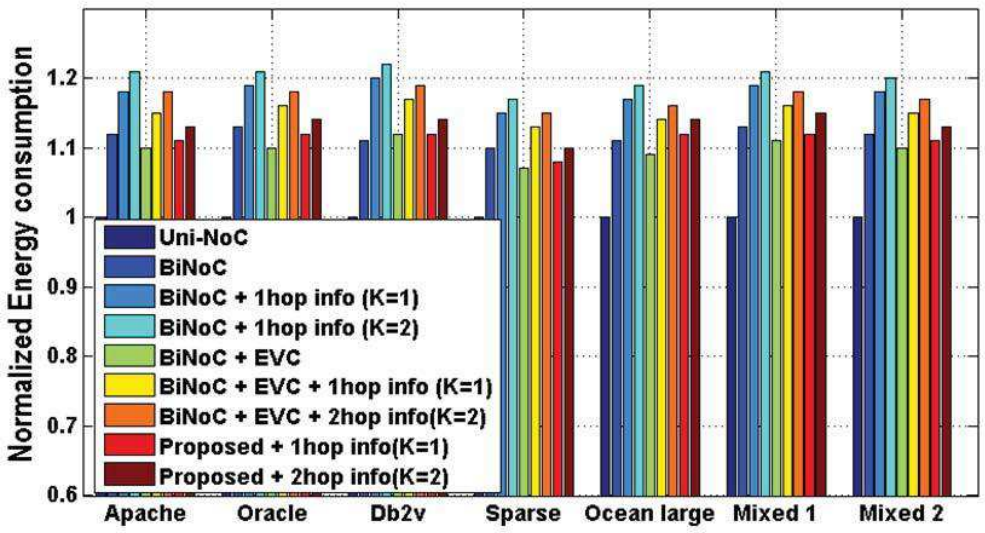

Figure 6.11: Energy comparison for benchmarks

to evaluate the channels status can reduce the average latency by $15\% - 60\%$. For example, after applying the fitness functions, the latency of the sparse benchmark is reduced to 42% compared with BiNoC routing. The static EVC approach further improves the performance by $6\% - 28\%$ compared with its BiNoC counterparts for both random selection or fitness function based selection strategies. As in the sparse application, the EVC+BiNoC+2 hop info routing improves the latency by 25.2% over the BiNoC + 2 hop info routing. If a central hub router is added into each region in our proposed NoC architecture, the performance can be further improved by $5\% - 10\%$. As shown in Figure 6.9, for the sparse application, the proposed NoC architecture with $K = 2$ speculative routing improves the latency by 6.5%.

Figure 6.10 shows the percentage of packets that utilize the EVCs and hub routers when running different benchmarks. As shown, on average, 5.5% of packets utilize the express hub paths to send packets to the destination directly, while 22.9% packets utilize EVCs at least once at run time.

Next, we evaluate the energy overhead of different NoC architectures while running these applications. For this purpose, we have modified the Orion power model [44] to consider the extra energy consumption over the unidirectional NoC. For the BiNoC, the additional energy consumption is mainly due to the larger crossbar size ($2N \times 2N$), and the channel direction control logic [101]. We modified the energy consumption of the crossbar in the Orion power model and added the

energy overhead for the link direction switching according to [101]. For BiNoC + EVC architecture, the additional energy consumption for the EVC latch and the EVC arbiter is discussed in [113] and we add this to the Orion model. For the energy consumption of the proposed architecture, the hub latches in the normal router are modeled similarly to the EVC, while the energy consumption of the hub router is modeled as a four-port, four-VC normal router as in the Orion power model. For the calculation of the fitness values for each channel, we need additional adders and multipliers to calculate the fitness values every 100 cycles. This energy overhead is modeled in a similar way as in [114].

In Figure 6.11, the energy consumption of the unidirectional NoC architecture is normalized to one. The BiNoC approach has about 10%-20% energy overhead due to a large crossbar size and the link direction switching overhead [101]. The EVC and regional hub approach can effectively reduce the energy by allowing some packets use the express paths and bypassing the router pipeline stages. In Figure 6.11, the energy overhead for the proposed NoC architecture with 2-step selection strategy is 13% on average.

6.5.5 Overhead evaluation

We have modified the Orion 2.0 area model [44] to evaluate the area overhead. In this evaluation, we target a 1.0GHz operating frequency under the 65nm technology. When evaluating the BiNoC, the bi-directional link direction control logic is considered as in [101]. As shown in Figure 6.12, the area numbers of source/sink nodes and bypass nodes in the BiNoC+EVC architecture, as well as the normal router and hub router in the proposed architecture are compared against the baseline BiNoC design which is normalized to one. Router (1,2) in Figure 6.3 is chosen when evaluating the normal routers in the proposed architecture, since this router is the most complicated one with an extra port connected to the hub router. In Figure 6.12, compared to BiNoC baseline router, the normal router of the proposed architecture has an overhead of about 13%. Together with the additional hub router in each 4×4 region, the proposed NoC architecture has an area overhead of 19%.

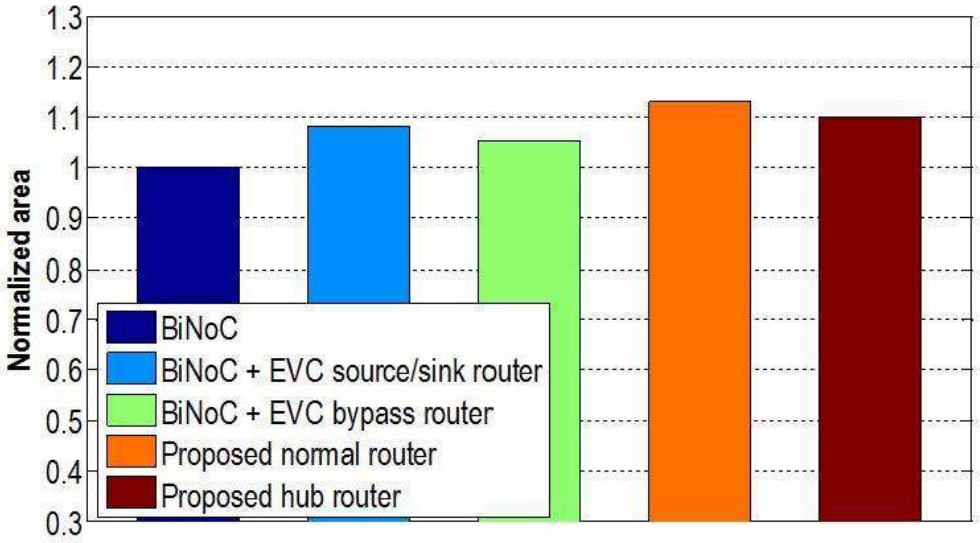

Figure 6.12: Area overhead comparison

6.6 Conclusion

In this chapter, we have proposed a flexible NoC architecture which utilizes bidirectional channels, EVC channels and regional hub routers to improve the NoC performance. We have also proposed a traffic-aware, adaptive routing algorithm that considers the characteristic of the bidirectional links. Simulation results on synthetic traffic, as well as real world benchmarks, show that the proposed NoC architecture, together with the routing algorithm, can significantly reduce the latency by as much as 80% with a small energy and area overhead.

CHAPTER 7

CONCLUSIONS AND FUTURE WORK

7.1 Research summary

The work discussed in this thesis focus on the issues of achieving a high performance NoC design for future multicore systems. In particular, starting from the offline design space exploration, we have proposed a more accurate NoC model to predict the latency performance under various routing algorithm and placement choices. We first identified the limitations of conventional simulations and queuing-theory-based performance evaluation methods. Motivated by combining the high accuracy in simulations and fast speed in analytical models, we have applied machine learning techniques to combine the advantages of simulations and mathematical formalism. In Chapter 2, we have presented a support vector regression based latency model for evaluating NoC performance in the synthesis inner loop. Through learning from the typical training data, the SVR-NoC achieves better accuracy and similar speedup performance compared to the queuing models, which will benefit the exploration of the design space offline.

Then, in Chapter 3 and Chapter 4, we have explored the NoC routing algorithm designs for different objectives and constraints. Specifically, we have discussed the routing algorithms for the thermal-awareness and fault-tolerance purposes. In Chapter 3, in order to reduce the hotspot temperature of the whole system while maintaining the latency and throughput performance, we have proposed an application-specific path set finding algorithm which has a high adaptivity to distribute traffic and ensures deadlock-free property. Then, a linear programming problem is formulated to compute the optimal ratio of using the paths offline. The router architecture which supports the proposed ratio-based adaptive routing is also discussed. In Chapter 4, we have proposed an adaptive routing algorithm to tackle run time router faults. In this chapter, we proposed a dynamic buffer swapping algorithm and a dynamic crossbar Mux swapping algorithm to maximally maintain the network connectivity under the buffer and crossbar faults, respectively. Higher packet acceptance rate and lower latency can be achieved

compared to previous re-routing based methods.

In Chapter 5 and Chapter 6, we have proceeded to explore the architectures for optimizing the system performance. We have proposed our new NoC design based on the usage of the emerging bi-directional channels. More precisely, we have designed two new flexible NoC architectures using bi-directional links to further improve the throughput and latency performance. In Chapter 5, the overall NoC topology is unchanged. We focused on the router design to improve the usage of bi-directional channel bandwidth. Towards this end, we have proposed a flit-level speedup scheme (FSNoC) to allow two flits from the same packet to be transmitted simultaneously. In this way, higher bandwidth utilization and throughput can be achieved compared to the conventional NoCs equipped with bi-directional channels. In Chapter 6, we have proposed a new NoC topology which takes the the advantage of several express paths provided in the system: 1) bi-directional channels, 2) express paths and 3) regional hub router based shortcut paths. An adaptive routing algorithm is designed to dynamically choose among these paths. We have demonstrated the effectiveness of the proposed new architecture using both synthetic traffic and real multicore workloads.

7.2 Future research directions

Additional work is needed to extend this thesis work. In future, we expect to conduct the researches in the following directions:

1) Exploring the current learning based NoC models: We will explore and compare several other learning methods besides support vector regression (*e.g.,* neural networks, Gaussian process). We will compare the model accuracy with the SVR-based model. Also, we will explore the inclusion of different features to improve the learning accuracy and reduce the training data size as well as learning time. We plan to study the importance of each features in the feature set and develop the criteria of how to choose features in the training stage. Moreover, we will develop a more flexible platform to extract arbitrary training data from the simulation.

2) Applying learning techniques to model the worst-case delay: Besides the average latency modeling in this work, we will also apply the learning techniques to predict the worst-case delays for the NoCs with real-time deadline

requirement. We will explore the effects of different learning methods on the worst case delay bound tightness.

3) Hardware implementation for the flexible NoC architecture: For our proposed NoC architecture, we plan to implement a FPGA-based prototype and carry out a thorough comparison and analysis on the power, area overhead.

4) More applications based on the proposed NoC platform: In future, we also expect to see more applications which is implemented based on the proposed NoC architectures. Specifically, we have started to apply the proposed NoC platform for a specific biological application, *i.e.,* the protein folding calculation and prediction.

5) Hardware implementations and detail evaluations of FSNoC under the existence of long wire links: As the bi-directional repeaters introduced in Chapter 5 may increase the critical path length and power overhead, a more accurate hardware overhead evaluation of FSNoC should also consider the physical VLSI implementation and optimization of the bi-repeaters.

6) Run-time thermal aware routing for reducing the hotspot temperature: The proposed thermal aware routing is based on application task graph, which only characterizes the average-case communication bandwidth requirement. Therefore, the routing is done in the offline phase and cannot capture the burst traffic arrivals at run time. In order to consider the fluctuation of power/thermal profile at run time, a fine-grained on-line routing algorithm is needed which determines the routing path dynamically based on the current temperature of the routers and PEs.

REFERENCES

[1] "Transistor count statistics," http://www.wikipedia.org/wiki/Transistor_count, 2013.

[2] John L. Hennessy and David A. Patterson, *Computer Architecture, Fourth Edition: A Quantitative Approach*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2012.

[3] Jan M. Rabaey, *Digital Integrated Circuits: A Design Perspective*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2008.

[4] "Design of energy-efficient on-chip networks," http://www.rle.mit.edu/isg/documents/2010_NOC_tutorial_vladimir.pdf, 2010.

[5] Maurizio Palesi and Masoud Daneshtalab (Eds.), *Routing Algorithms in Networks-on-Chip*, Springer, 2014.

[6] N. E. Jerger and L. S. Peh, *On-chip Networks*, Morgan, New York, 2009.

[7] Y. Wu, G. Min, M. Ould-Khaoua, H. Yin, and L. Wang, "Analytical modelling of networks in multicomputer systems under bursty and batch arrival traffic," *The Journal of Supercomputing*, vol. 51, no. 2, pp. 115–130, 2010.

[8] C.M. Bishop, *Pattern recognition and machine learning*, Springer, 2006.

[9] Shu-Yen Lin, Wen-Chung Shen, Chan-Cheng Hsu, Chih-Hao Chao, and An-Yeu Wu, "Fault-tolerant router with built-in self-test/self-diagnosis and fault-isolation circuits for 2d-mesh based chip multiprocessor systems," in *VLSI Design, Automation and Test, 2009. VLSI-DAT '09. International Symposium on*, 2009, pp. 72–75.

[10] A. Kohler, G. Schley, and M. Radetzki, "Fault tolerant network on chip switching with graceful performance degradation," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 29, no. 6, pp. 883–896, 2010.

[11] Robert R. Schaller, "Moore's law: Past, present, and future," *IEEE Spectr.*, vol. 34, no. 6, pp. 52–59, June 1997.

[12] "Intel xeon processors," `http://www.wikipedia.org/wiki/Xeon`, 2013.

[13] "Nvidia geforce gtx780," `http://www.techpowerup.com/reviews/NVIDIA/GeForce_GTX_78` 2013.

[14] Lap Fai Leung, *Designing high-performance and low-energy real-time embedded systems based on single-core and multi-cores structures (Ph.D. Thesis)*, HKUST, Hong Kong, 2007.

[15] F. Clermidy, C. Bernard, R. Lemaire, J. Martin, I. Miro-Panades, Y. Thonnart, P. Vivet, and N. Wehn, "A 477mw noc-based digital baseband for mimo 4g sdr," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, 2010, pp. 278–279.

[16] S. R. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, A. Singh, T. Jacob, S. Jain, V. Erraguntla, C. Roberts, Y. Hoskote, N. Borkar, and S. Borkar, "An 80-tile sub-100-w teraflops processor in 65-nm cmos," *Solid-State Circuits, IEEE Journal of*, vol. 43, no. 1, pp. 29–41, 2008.

[17] Yaoyao Ye, Jiang Xu, Xiaowen Wu, Wei Zhang, Xuan Wang, M. Nikdast, Zhehui Wang, and Weichen Liu, "System-level modeling and analysis of thermal effects in optical networks-on-chip," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 21, no. 2, pp. 292–305, Feb 2013.

[18] S. Poddar, P. Ghosal, P. Mukherjee, S. Samui, and H. Rahaman, "Design of an noc with on-chip photonic interconnects using adaptive cdma links," in *SOC Conference (SOCC), 2012 IEEE International*, Sept 2012, pp. 352–357.

[19] D. Fick, A. DeOrio, Jin Hu, V. Bertacco, D. Blaauw, and D Sylvester, "Vicis: A reliable network for unreliable silicon," in *Design Automation Conference, 2009. DAC '09. 46th ACM/IEEE*, 2009, pp. 812–817.

[20] William Dally and Brian Towles, *Principles and Practices of Interconnection Networks*, Morgan Kaufmann, 2003.

[21] Y. Thonnart, P. Vivet, and F. Clermidy, "A fully-asynchronous low-power framework for gals noc integration," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2010*, 2010, pp. 33–38.

[22] "A comparison of network-on-chip and busses," `http://www.ict.kth.se/courses/IL2207/0810/docs/noc_whitepaper.pdf`, 2008.

[23] M. Palesi, R. Holsmark, S. Kumar, and V. Catania, "Application specific routing algorithms for networks on chip," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 20, no. 3, pp. 316–330, 2009.

[24] A. E. Kiasari, Z. Lu, and A. Jantsch, "An analytical latency model for networks-on-chip," *Very Large Scale Integration (VLSI) Systems, IEEE Trans. on*, vol. PP, no. 99, pp. 1 –11, 2012.

[25] M. Mirza-Aghatabar, S. Koohi, S. Hessabi, and M. Pedram, "An empirical investigation of mesh and torus noc topologies under different routing algorithms and traffic models," in *Digital System Design Architectures, Methods and Tools, 2007. DSD 2007. 10th Euromicro Conference on*, 2007, pp. 19–26.

[26] J. Kim, J. Balfour, and W.J. Dally, "Flattened butterfly topology for on-chip networks," in *Microarchitecture, 2007. MICRO 2007. 40th Annual IEEE/ACM International Symposium on*, 2007, pp. 172–182.

[27] M. Ahmed and R. Kumar, "Parameterized path-based, randomized, oblivious, minimal routing in 3d mesh noc," in *TENCON 2012 - 2012 IEEE Region 10 Conference*, 2012, pp. 1–6.

[28] V. Dvorak and J. Jaros, "Optimizing collective communications on 2d-mesh and fat tree noc," in *Networks (ICN), 2010 Ninth International Conference on*, 2010, pp. 22–27.

[29] M.B. Taylor, J. Kim, J. Miller, D. Wentzlaff, F. Ghodrat, B. Greenwald, H. Hoffman, P. Johnson, Jae-Wook Lee, W. Lee, A. Ma, A. Saraf, M. Seneski, N. Shnidman, V. Strumpen, M. Frank, S. Amarasinghe, and A. Agarwal, "The raw microprocessor: a computational fabric for software

circuits and general-purpose programs," *Micro, IEEE*, vol. 22, no. 2, pp. 25–35, 2002.

[30] Jose Duato, Sudhakar Yalamanchili, and Ni Lionel, *Interconnection Networks: An Engineering Approach*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002.

[31] Parviz Kermani and Leonard Kleinrock, "Virtual cut-through: a new computer communication switching technique," *Computer Networks*, vol. 3, pp. 267–286, 1979.

[32] Ling Xin and Chiu-Sing Choy, "A low-latency noc router with lookahead bypass," in *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, May 2010, pp. 3981–3984.

[33] Mingche Lai, Lei Gao, Nong Xiao, and Zhiying Wang, "An accurate and efficient performance analysis approach based on queuing model for network on chip," in *proc. ICCAD 2009.*, nov. 2009, pp. 563 –570.

[34] Z. Guz, I. Walter, E. Bolotin, I. Cidon, R. Ginosar, and A. Kolodny, "Network delays and link capacities in application-specific wormhole nocs," 2007.

[35] U.Y. Ogras, P. Bogdan, and R. Marculescu, "An analytical approach for network-on-chip performance analysis," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Trans. on*, vol. 29, no. 12, pp. 2001 –2013, dec. 2010.

[36] Y. Ben-Itzhak, I. Cidon, and A. Kolodny, "Delay analysis of wormhole based heterogeneous noc," in *Networks on Chip (NoCS), 2011 Fifth IEEE/ACM International Symposium on*, 2011, pp. 161–168.

[37] P. Bogdan and R. Marculescu, "Non-stationary traffic analysis and its implications on multicore platform design," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 30, no. 4, pp. 508–519, 2011.

[38] V. Vapnik, *Statistical Learning theory*, John Wiley and Sons, 1998.

[39] O. Lysne, "Towards a generic analytical model of wormhole routing networks," *microprocessors and microsystems*, vol. 21, no. 7-8, pp. 491–498, 1998.

[40] M. Bakhouya, S. Suboh, J. Gaber, and T. El-Ghazawi, "Analytical modeling and evaluation of on-chip interconnects using network calculus," in *proc. NoCS 2009.*, may 2009, pp. 74 –79.

[41] N. Nikitin and J. Cortadella, "A performance analytical model for network-on-chip with constant service time routers," in *proc. ICCAD 2009.*, nov. 2009, pp. 571 –578.

[42] Po-Chi Hu and Leonard Kleinrock, "An analytical model for wormhole routing with finite size input buffers," 1997.

[43] A.B. Kahng, Bill Lin, and K. Samadi, "Improved on-chip router analytical power and area modeling," in *Design Automation Conference (ASP-DAC), 2010 15th Asia and South Pacific*, 2010, pp. 241–246.

[44] A.B. Kahng, Bin Li, Li-Shiuan Peh, and K. Samadi, "Orion 2.0: A power-area simulator for interconnection networks," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 20, no. 1, pp. 191–196, 2012.

[45] Andrew B. Kahng, Bill Lin, and Siddhartha Nath, "Explicit modeling of control and data for improved noc router estimation," in *Proceedings of the 49th Annual Design Automation Conference*, New York, NY, USA, 2012, DAC '12, pp. 392–397, ACM.

[46] Da-Cheng Juan and Diana Marculescu, "Power-aware performance increase via core/uncore reinforcement control for chip-multiprocessors," in *Proceedings of the 2012 ACM/IEEE International Symposium on Low Power Electronics and Design*, New York, NY, USA, 2012, ISLPED '12, pp. 97–102, ACM.

[47] M. Ebrahimi, M. Daneshtalab, F. Farahnakian, J. Plosila, P. Liljeberg, M. Palesi, and H. Tenhunen, "Haraq: Congestion-aware learning model for highly adaptive routing algorithm in on-chip networks," in *Networks on Chip (NoCS), 2012 Sixth IEEE/ACM International Symposium on*, 2012, pp. 19–26.

[48] Terrence Mak, Peter Y.K. Cheung, Wayne Luk, and Kai Pui Lam, "A dp-network for optimal dynamic routing in network-on-chip," in *Proceedings of the 7th IEEE/ACM International Conference on Hardware/Software Codesign and System Synthesis*, New York, NY, USA, 2009, CODES+ISSS '09, pp. 119–128, ACM.

[49] D. D. Kouvatsos, A. SALAM, and M. Ould-Khaoua, "Performance modeling of wormhole-routed hypercubes with bursty traffic and finite buffers," *Int J Simul Pract Syst Sci Technol*, vol. 6, pp. 69–81, 2005.

[50] M. Arjomand and H. Sarbazi-Azad, "A comprehensive power-performance model for nocs with multi-flit channel buffers," in *Proceedings of the 23rd international conference on Supercomputing*, New York, NY, USA, 2009, ICS '09, pp. 470–478, ACM.

[51] G. Min and M. Ould-Khaoua, "A performance model for wormhole-switched interconnection networks under self-similar traffic," *Computers, IEEE Transactions on*, vol. 53, no. 5, pp. 601–613, 2004.

[52] E. Fischer and G.P. Fettweis, "An accurate and scalable analytic model for round-robin arbitration in network-on-chip," in *Networks on Chip (NoCS), 2013 Seventh IEEE/ACM International Symposium on*, 2013, pp. 1–8.

[53] G. Casale, E. Z. Zhang, and E. Smirni, "Kpc-toolbox: Simple yet effective trace fitting using markovian arrival processes," in *Quantitative Evaluation of Systems, 2008. QEST '08. Fifth International Conference on*, 2008, pp. 83–92.

[54] M. Ould-Khaoua, "A performance model for duato's fully adaptive routing algorithm in k-ary n-cubes," *Computers, IEEE Transactions on*, vol. 48, no. 12, pp. 1297–1304, 1999.

[55] W. Dally, "Virtual-channel flow control," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 3, no. 2, pp. 194–205, 1992.

[56] A. E. Kiasari, D. Rahmati, H. Sarbazi-Azad, and S. Hessabi, "A markovian performance model for networks-on-chip," in *Parallel, Distributed and Network-Based Processing, 2008. PDP 2008. 16th Euromicro Conference on*, 2008, pp. 157–164.

[57] M. A. El-Affendi and D. D. Kouvatsos, "A maximum entropy analysis of the m/g/1 and g/m/1 queueing systems at equilibrium," *Acta Informatica*, vol. 19, no. 4, pp. 339–355, 1983.

[58] Zhiliang Qian, Da-Cheng Juan, Paul Bogdan, Chi-Ying Tsui, Diana Marculescu, and Radu Marculescu, "Svr-noc: A performance analysis tool for network-on-chips using learning-based support vector regression model," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2013*, March 2013, pp. 354–357.

[59] *libsvm,http://www.csie.ntu.edu.tw/ cjlin/libsvm*, 2012.

[60] *LS-SVMlab,http://www.esat.kuleuven.be/sista/lssvmlab*, 2012.

[61] *Booksim 2.0, http://nocs.stanford.edu/booksim.html*, 2012.

[62] A. Pullini, F. Angiolini, P. Meloni, D. Atienza, S. Murali, L. Raffo, G. De Micheli, and L. Benini, "Noc design and implementation in 65nm technology," in *Networks-on-Chip, 2007. NOCS 2007. First International Symposium on*, 2007, pp. 273–282.

[63] L. Shang, Li-Shiuan Peh, A Kumar, and N.K. Jha, "Temperature-aware on-chip networks," *Micro, IEEE*, vol. 26, no. 1, pp. 130–139, 2006.

[64] B.C. Schafer and Taewhan Kim, "Hotspots elimination and temperature flattening in vlsi circuits," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 16, no. 11, pp. 1475–1487, 2008.

[65] W. Hung, C. Addo-Quaye, T. Theocharides, Y. Xie, N. Vijakrishnan, and M.J. Irwin, "Thermal-aware ip virtualization and placement for networks-on-chip architecture," in *Computer Design: VLSI in Computers and Processors, 2004. ICCD 2004. Proceedings. IEEE International Conference on*, 2004, pp. 430–437.

[66] S. Vangal, A. Singh, J. Howard, S. Dighe, N. Borkar, and A. Alvandpour, "A 5.1ghz 0.34mm2 router for network-on-chip applications," in *VLSI Circuits, 2007 IEEE Symposium on*, 2007, pp. 42–43.

[67] Jingcao Hu and R. Marculescu, "Energy- and performance-aware mapping for regular noc architectures," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 24, no. 4, pp. 551 – 562, april 2005.

[68] D. Bertozzi, A. Jalabert, S. Murali, R. Tamhankar, S. Stergiou, L. Benini, and G. De Micheli, "Noc synthesis flow for customized domain specific multiprocessor systems-on-chip," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 16, no. 2, pp. 113–129, 2005.

[69] M. Daneshtalab, A. Sobhani, A. Afzali-Kusha, O. Fatemi, and Z. Navabi, "Noc hot spot minimization using antnet dynamic routing algorithm," in *Application-specific Systems, Architectures and Processors, 2006. ASAP '06. International Conference on*, 2006, pp. 33–38.

[70] G.M. Link and N. Vijaykrishnan, "Hotspot prevention through runtime reconfiguration in network-on-chip," in *Design, Automation and Test in Europe, 2005. Proceedings*, 2005, pp. 648–649 Vol. 1.

[71] D. Atienza and E. Martinez, "Inducing thermal-awareness in multicore systems using networks-on-chip," in *VLSI, 2009. ISVLSI '09. IEEE Computer Society Annual Symposium on*, 2009, pp. 187–192.

[72] Ying-Cherng Lan, M.C. Chen, A.P. Su, Yu-Hen Hu, and Sao-Jie Chen, "Fluidity concept for noc: A congestion avoidance and relief routing scheme," in *SOC Conference, 2008 IEEE International*, 2008, pp. 65–70.

[73] Chih-Hao Chao, Kun-Chih Chen, and An-Yeu Wu, "Routing-based traffic migration and buffer allocation schemes for 3-d network-on-chip systems with thermal limit," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 21, no. 11, pp. 2118–2131, 2013.

[74] Zhiliang Qian and Chi-Ying Tsui, "A thermal-aware application specific routing algorithm for network-on-chip design," in *Design Automation Conference (ASP-DAC), 2011 16th Asia and South Pacific*, Jan 2011, pp. 449–454.

[75] J. Duato, "A necessary and sufficient condition for deadlock-free adaptive routing in wormhole networks," *Parallel and Distributed Systems, IEEE*

Transactions on, vol. 6, no. 10, pp. 1055–1067, 1995.

[76] Robert Tarjan, "Depth first search and linear graph algorithms," *SIAM Journal on Computing*, 1972.

[77] "Hotspot 5.0 temperature modeling tool," `http://lava.cs.virginia.edu/HotSpot`, 2008.

[78] "Cvx: Matlab software for disciplined convex programming," `http://cvxr.com/cvx`, 2010.

[79] "Noxim simulator," `http://www.noxim.org`, 2010.

[80] "Nigram simulator," `http://nigram.ecs.soton.ac.uk`.

[81] D. Bertozzi and L. Benini, "Xpipes: a network-on-chip architecture for gigascale systems-on-chip," *Circuits and Systems Magazine, IEEE*, vol. 4, no. 2, pp. 18–31, 2004.

[82] M. Palesi, S. Kumar, and V. Catania, "Bandwidth-aware routing algorithms for networks-on-chip platforms," *Computers Digital Techniques, IET*, vol. 3, no. 5, pp. 413–429, 2009.

[83] Christopher J. Glass and Lionel M. Ni, "The turn model for adaptive routing," *SIGARCH Comput. Archit. News*, vol. 20, no. 2, pp. 278–287, Apr. 1992.

[84] Rajendra V. Boppana and Suresh Chalasani, "A comparison of adaptive wormhole routing algorithms," in *Proceedings of the 20th Annual International Symposium on Computer Architecture*, New York, NY, USA, 1993, ISCA '93, pp. 351–360, ACM.

[85] Zhen Zhang, A. Greiner, and S. Taktak, "A reconfigurable routing algorithm for a fault-tolerant 2d-mesh network-on-chip," in *Design Automation Conference, 2008. DAC 2008. 45th ACM/IEEE*, 2008, pp. 441–446.

[86] D. Fick, A. DeOrio, G. Chen, V. Bertacco, D Sylvester, and D. Blaauw, "A highly resilient routing algorithm for fault-tolerant nocs," in *Design, Automation Test in Europe Conference Exhibition, 2009. DATE '09.*, 2009, pp. 21–26.

[87] A. Hosseini, T. Ragheb, and Y. Massoud, "A fault-aware dynamic routing algorithm for on-chip networks," in *Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on*, 2008, pp. 2653–2656.

[88] Dongkook Park, C. Nicopoulos, Jongman Kim, N. Vijaykrishnan, and C.R. Das, "Exploring fault-tolerant network-on-chip architectures," in *Dependable Systems and Networks, 2006. DSN 2006. International Conference on*, 2006, pp. 93–104.

[89] Zhiliang Qian, Ying-Fei Teh, and Chi-Ying Tsui, "A fault-tolerant network-on-chip design using dynamic reconfiguration of partial-faulty routing resources," in *VLSI and System-on-Chip (VLSI-SoC), 2011 IEEE/IFIP 19th International Conference on*, Oct 2011, pp. 192–195.

[90] D. Frazzetta, G. Dimartino, M. Palesi, S. Kumar, and V. Catania, "Efficient application specific routing algorithms for noc systems utilizing partially faulty links," in *Digital System Design Architectures, Methods and Tools, 2008. DSD '08. 11th EUROMICRO Conference on*, 2008, pp. 18–25.

[91] Ge-Ming Chiu, "The odd-even turn model for adaptive routing," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 11, no. 7, pp. 729–738, 2000.

[92] J. Duato, "A new theory of deadlock-free adaptive routing in wormhole networks," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 4, no. 12, pp. 1320–1331, 1993.

[93] K.V. Anjan and T.M. Pinkston, "An efficient, fully adaptive deadlock recovery scheme: Disha," in *Computer Architecture, 1995. Proceedings., 22nd Annual International Symposium on*, 1995, pp. 201–210.

[94] M. Palesi, G. Longo, S. Signorino, R. Holsmark, S. Kumar, and V. Catania, "Design of bandwidth aware and congestion avoiding efficient routing algorithms for networks-on-chip platforms," in *Networks-on-Chip, 2008. NoCS 2008. Second ACM/IEEE International Symposium on*, 2008, pp. 97–106.

[95] L. Benini and G. De Micheli, "Networks on chips: a new soc paradigm," *Computer*, vol. 35, no. 1, pp. 70 –78, jan 2002.

[96] C.A. Nicopoulos, Dongkook Park, Jongman Kim, N. Vijaykrishnan, M.S. Yousif, and C.R. Das, "Vichar: A dynamic virtual channel regulator for network-on-chip routers," in *Microarchitecture, 2006. MICRO-39. 39th Annual IEEE/ACM International Symposium on*, 2006, pp. 333–346.

[97] Jie Meng, Chao Chen, Ayse Kivilcim Coskun, and Ajay Joshi, "Run-time energy management of manycore systems through reconfigurable interconnects," in *Proceedings of the 21st edition of the great lakes symposium on Great lakes symposium on VLSI*, New York, NY, USA, 2011, GLSVLSI '11, pp. 43–48, ACM.

[98] Lei Wang, Poornachandran Kumar, Ki Hwan Yum, and Eun Jung Kim, "Apcr: an adaptive physical channel regulator for on-chip interconnects," in *Proceedings of the 21st international conference on Parallel architectures and compilation techniques*. 2012, PACT '12, pp. 87–96, ACM.

[99] Asit K. Mishra, N. Vijaykrishnan, and Chita R. Das, "A case for heterogeneous on-chip interconnects for cmps," in *Proceedings of the 38th annual international symposium on Computer architecture*, New York, NY, USA, 2011, ISCA '11, pp. 389–400, ACM.

[100] M.A. Al Faruque, T. Ebi, and J. Henkel, "Configurable links for runtime adaptive on-chip communication," in *Design, Automation Test in Europe Conference Exhibition, 2009. DATE '09.*, april 2009, pp. 256 –261.

[101] Ying-Cherng Lan, Hsiao-An Lin, Shih-Hsin Lo, Yu Hen Hu, and Sao-Jie Chen, "A bidirectional noc (binoc) architecture with dynamic self-reconfigurable channel," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 30, no. 3, pp. 427 –440, march 2011.

[102] Myong Hyon Cho, M. Lis, Keun Sup Shim, M. Kinsy, T. Wen, and S Devadas, "Oblivious routing in on-chip bandwidth-adaptive networks," in *Parallel Architectures and Compilation Techniques, 2009. PACT '09. 18th International Conference on*, 2009, pp. 181–190.

[103] Wen-Chung Tsai, Deng-Yuan Zheng, Sao-Jie Chen, and Yu-Hen Hu, "A fault-tolerant noc scheme using bidirectional channel," in *Design Automa-*

tion Conference (DAC), 2011 48th ACM/EDAC/IEEE, june 2011, pp. 918 –923.

[104] Shih-Hsin Lo, Ying-Cherng Lan, Hsin-Hsien Yeh, Wen-Chung Tsai, Yu-Hen Hu, and Sao-Jie Chen, "Qos aware binoc architecture," in *Parallel Distributed Processing (IPDPS), 2010 IEEE International Symposium on*, april 2010, pp. 1 –10.

[105] R. Hesse, J. Nicholls, and N.E. Jerger, "Fine-grained bandwidth adaptivity in networks-on-chip using bidirectional channels," in *Networks on Chip (NoCS), 2012 Sixth IEEE/ACM International Symposium on*, 2012, pp. 132–141.

[106] R. Mullins, A. West, and S. Moore, "Low-latency virtual-channel routers for on-chip networks," in *Computer Architecture, 2004. Proceedings. 31st Annual International Symposium on*, june 2004, pp. 188 – 197.

[107] A.K. Lusala, P. Manet, B. Rousseau, and J. Legat, "Noc implementation in fpga using torus topology," in *Field Programmable Logic and Applications, 2007. FPL 2007. International Conference on*, 2007, pp. 778–781.

[108] U.Y. Ogras and R. Marculescu, ""it's a small world after all": Noc performance optimization via long-range link insertion," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 14, no. 7, pp. 693–706, 2006.

[109] A. Morgenshtein, I. Cidon, A. Kolodny, and R. Ginosar, "Low-leakage repeaters for noc interconnects," in *Circuits and Systems, 2005. ISCAS 2005. IEEE International Symposium on*, 2005, pp. 600–603 Vol. 1.

[110] *E3S benchmarks, http://ziyang.eecs.umich.edu/dickrp/e3s/.*

[111] *Spec benchmarks, http://www.spec.org/web99*, 1999.

[112] Sheng Ma, N.E. Jerger, and Zhiying Wang, "Dbar: An efficient routing algorithm to support multiple concurrent applications in networks-on-chip," in *Computer Architecture (ISCA), 2011 38th Annual International Symposium on*, 2011, pp. 413–424.

[113] A Kumar, Li-Shiuan Peh, P. Kundu, and N.K. Jha, "Toward ideal on-chip communication using express virtual channels," *Micro, IEEE*, vol. 28, no. 1, pp. 80–90, 2008.

[114] G. Ascia, V. Catania, M. Palesi, and D. Patti, "Neighbors-on-path: A new selection strategy for on-chip networks," in *Embedded Systems for Real Time Multimedia, Proceedings of the 2006 IEEE/ACM/IFIP Workshop on*, 2006, pp. 79–84.

[115] P. Gratz, B. Grot, and S.W. Keckler, "Regional congestion awareness for load balance in networks-on-chip," in *High Performance Computer Architecture, 2008. HPCA 2008. IEEE 14th International Symposium on*, 2008, pp. 203–214.

[116] R.S. Ramanujam and Bill Lin, "Destination-based adaptive routing on 2d mesh networks," in *Architectures for Networking and Communications Systems (ANCS), 2010 ACM/IEEE Symposium on*, 2010, pp. 1–12.

[117] Zhiliang Qian, Paul Bogdan, Guopeng Wei, Chi-Ying Tsui, and Radu Marculescu, "A traffic-aware adaptive routing algorithm on a highly reconfigurable network-on-chip architecture," in *Proceedings of the Eighth IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis*, New York, NY, USA, 2012, CODES+ISSS '12, pp. 161–170, ACM.

[118] Chifeng Wang, Wen-Hsiang Hu, and N. Bagherzadeh, "Congestion-aware network-on-chip router architecture," in *Computer Architecture and Digital Systems (CADS), 2010 15th CSI International Symposium on*, 2010, pp. 137–144.

[119] B. Grot, J. Hestness, S.W. Keckler, and O. Mutlu, "Express cube topologies for on-chip interconnects," in *High Performance Computer Architecture, 2009. HPCA 2009. IEEE 15th International Symposium on*, 2009, pp. 163–174.

[120] Zhiliang Qian, Ying-Fei Teh, and Chi-Ying Tsui, "A flit-level speedup scheme for network-on-chips using self-reconfigurable bi-directional chan-

nels," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2012*, 2012, pp. 1295–1300.

[121] Minglong Shao, Anastassia Ailamaki, and Babak Falsafi, "Dbmbench: Fast and accurate database workload representation on modern microarchitecture," in *Proceedings of the 2005 Conference of the Centre for Advanced Studies on Collaborative Research*. 2005, CASCON '05, pp. 254–267, IBM Press.